


The RHIC & ATLAS Computing Facility Data Intensive Computing At The Scientific Frontier

Michael Ernst
BNL Physics Department
January 5, 2016


The Standard Model

A single, over-arching theory that includes all types of matter in the Universe and explains how they interact...




Must be self-consistent and explain all known physical observations!


The “Big Questions”




Magnitude = 1



Magnitude = 10^{25}




Magnitude = 10^{36}




Magnitude = 10^{38}

How does Gravity even fit in?



Where does this end?



How did we evolve from the Big Bang?



What is the nature of mass?



*And what is the more fundamental
underlying theory?*

Four Pillars Of Modern Experimental Particle Physics Research

Four classical pillars, resembling the Doric style, are arranged in a row. They are dark gray and have a subtle gradient. The pillars are positioned behind the text, with the text centered between the second and third pillars. The pillars have a fluted shaft and a capital with a volute.

Accelerators

Detectors

Theory

Computing



- 3.8 km circumference, in operation since 2000 (now last remaining collider in the U.S.)
- Multiple species of heavy ions up to 200 GeV/nucleon ($T \sim 4 \times 10^{12}$ K)
- Polarized protons up to 500 GeV c-o-m



The Large Hadron Collider at CERN

Lake Geneva

CMS


- Designed for 14 TeV p-p collisions at 40 MHz
- 10^{11} protons/bunch X 1500 bunches (10^9 int/sec)

LHCb


ALICE

ATLAS

CERN





Identifying and Measuring Particle



Why so big?

New Physics = High Mass
High Mass = High Momentum Tracks
High Momentum = Almost Straight
Big Field, Big Detector = Precision
Precision = Discovery!

The RHIC Detectors




LHC Point 1: The ATLAS Experiment

The ATLAS Collaboration


3000 Members

177 Institutes

38 Countries



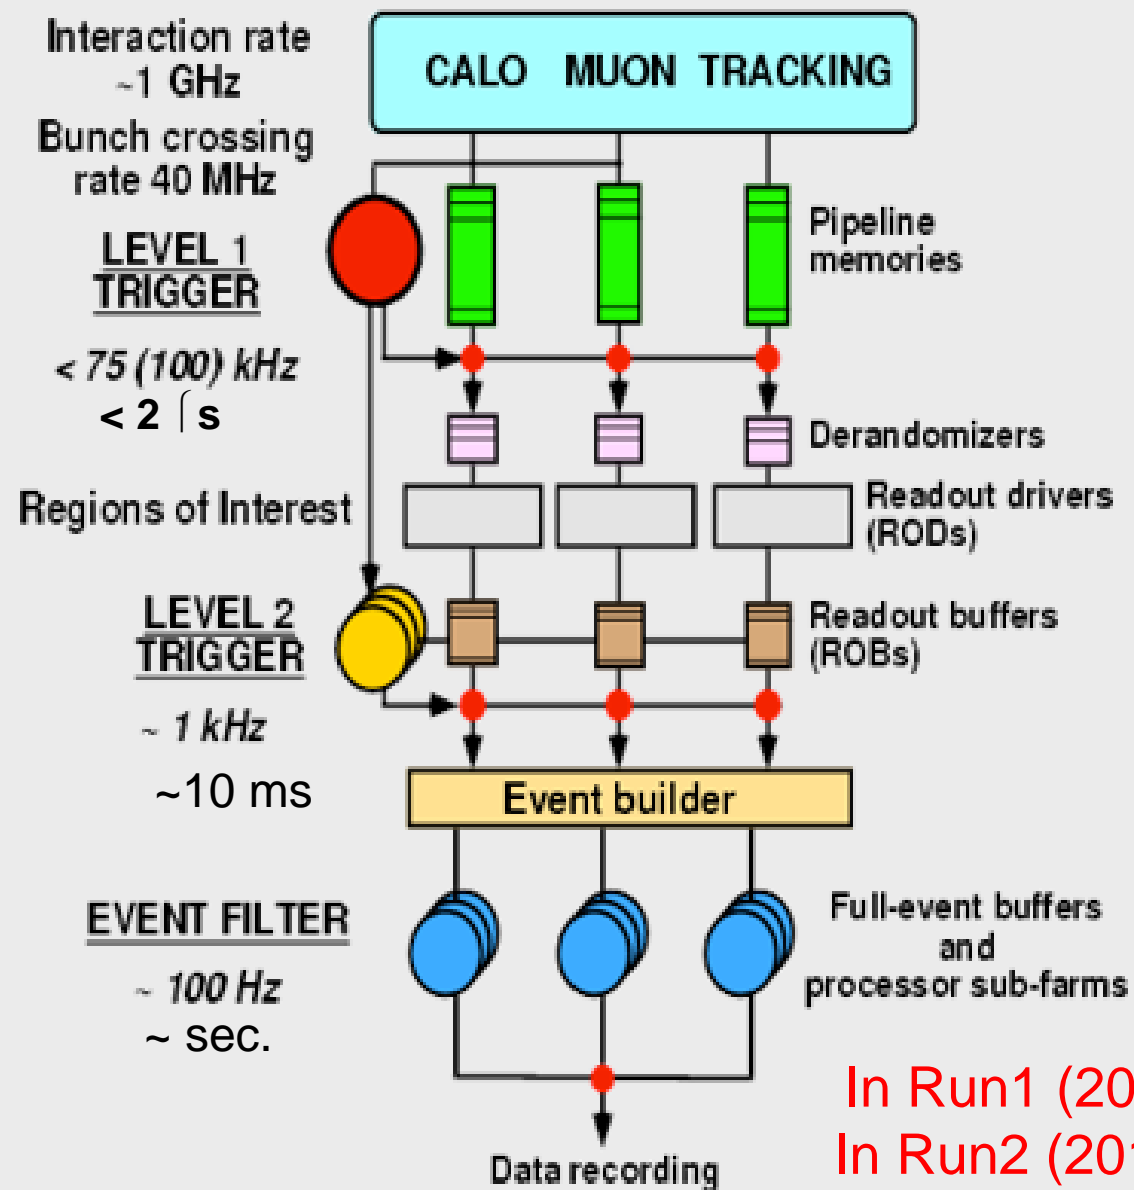
Electronic Channels At The ATLAS Detector



Length : ~ 46 m (150 ft)
Radius : ~ 12 m (40 ft)
Weight : ~ 7000 tons

~ 1800 miles of cables
~10 MW of electrical power

ATLAS Data Acquisition Rates



Physics selection of the 200 'best' out of ~1B interactions/sec:

40 MHz, 1 PB/sec

Level 1: Coarse calorimeter data and muon trigger chambers

75 kHz, 75 GB/sec

Level 2: Full information from all detectors in regions of interest

1 kHz, 1 GB/sec

Event Filter: Reconstruction of complete event using latest alignment and calibration data


In Run1 (2010-2012) 400 Hz, ~640 MB/sec

In Run2 (2015-2018) 1000Hz, ~1600 MB/s

62TB/day, 6 Petabyte/year of recorded raw data

LHC Beyond Run 2

To the High Luminosity LHC (HL-LHC)



Trigger rates, event complexity increase steadily through machine and detector upgrades

~15 PB/year LHC raw data in Run 1; ~130 PB/year in 2021 (all 4 LHC Experiments)

Very rough estimate for **new raw data per year in Run 4: 400 PB**

Raw data is only the beginning, e.g. ATLAS dataset now is ~160 PB, ~50% on disk

Event complexity: pile-up reaches ~150 at HL-LHC (now ~30), multiplicity up 8x

Plenty of challenges ahead!

DAQ/HLT Upgrade Plans

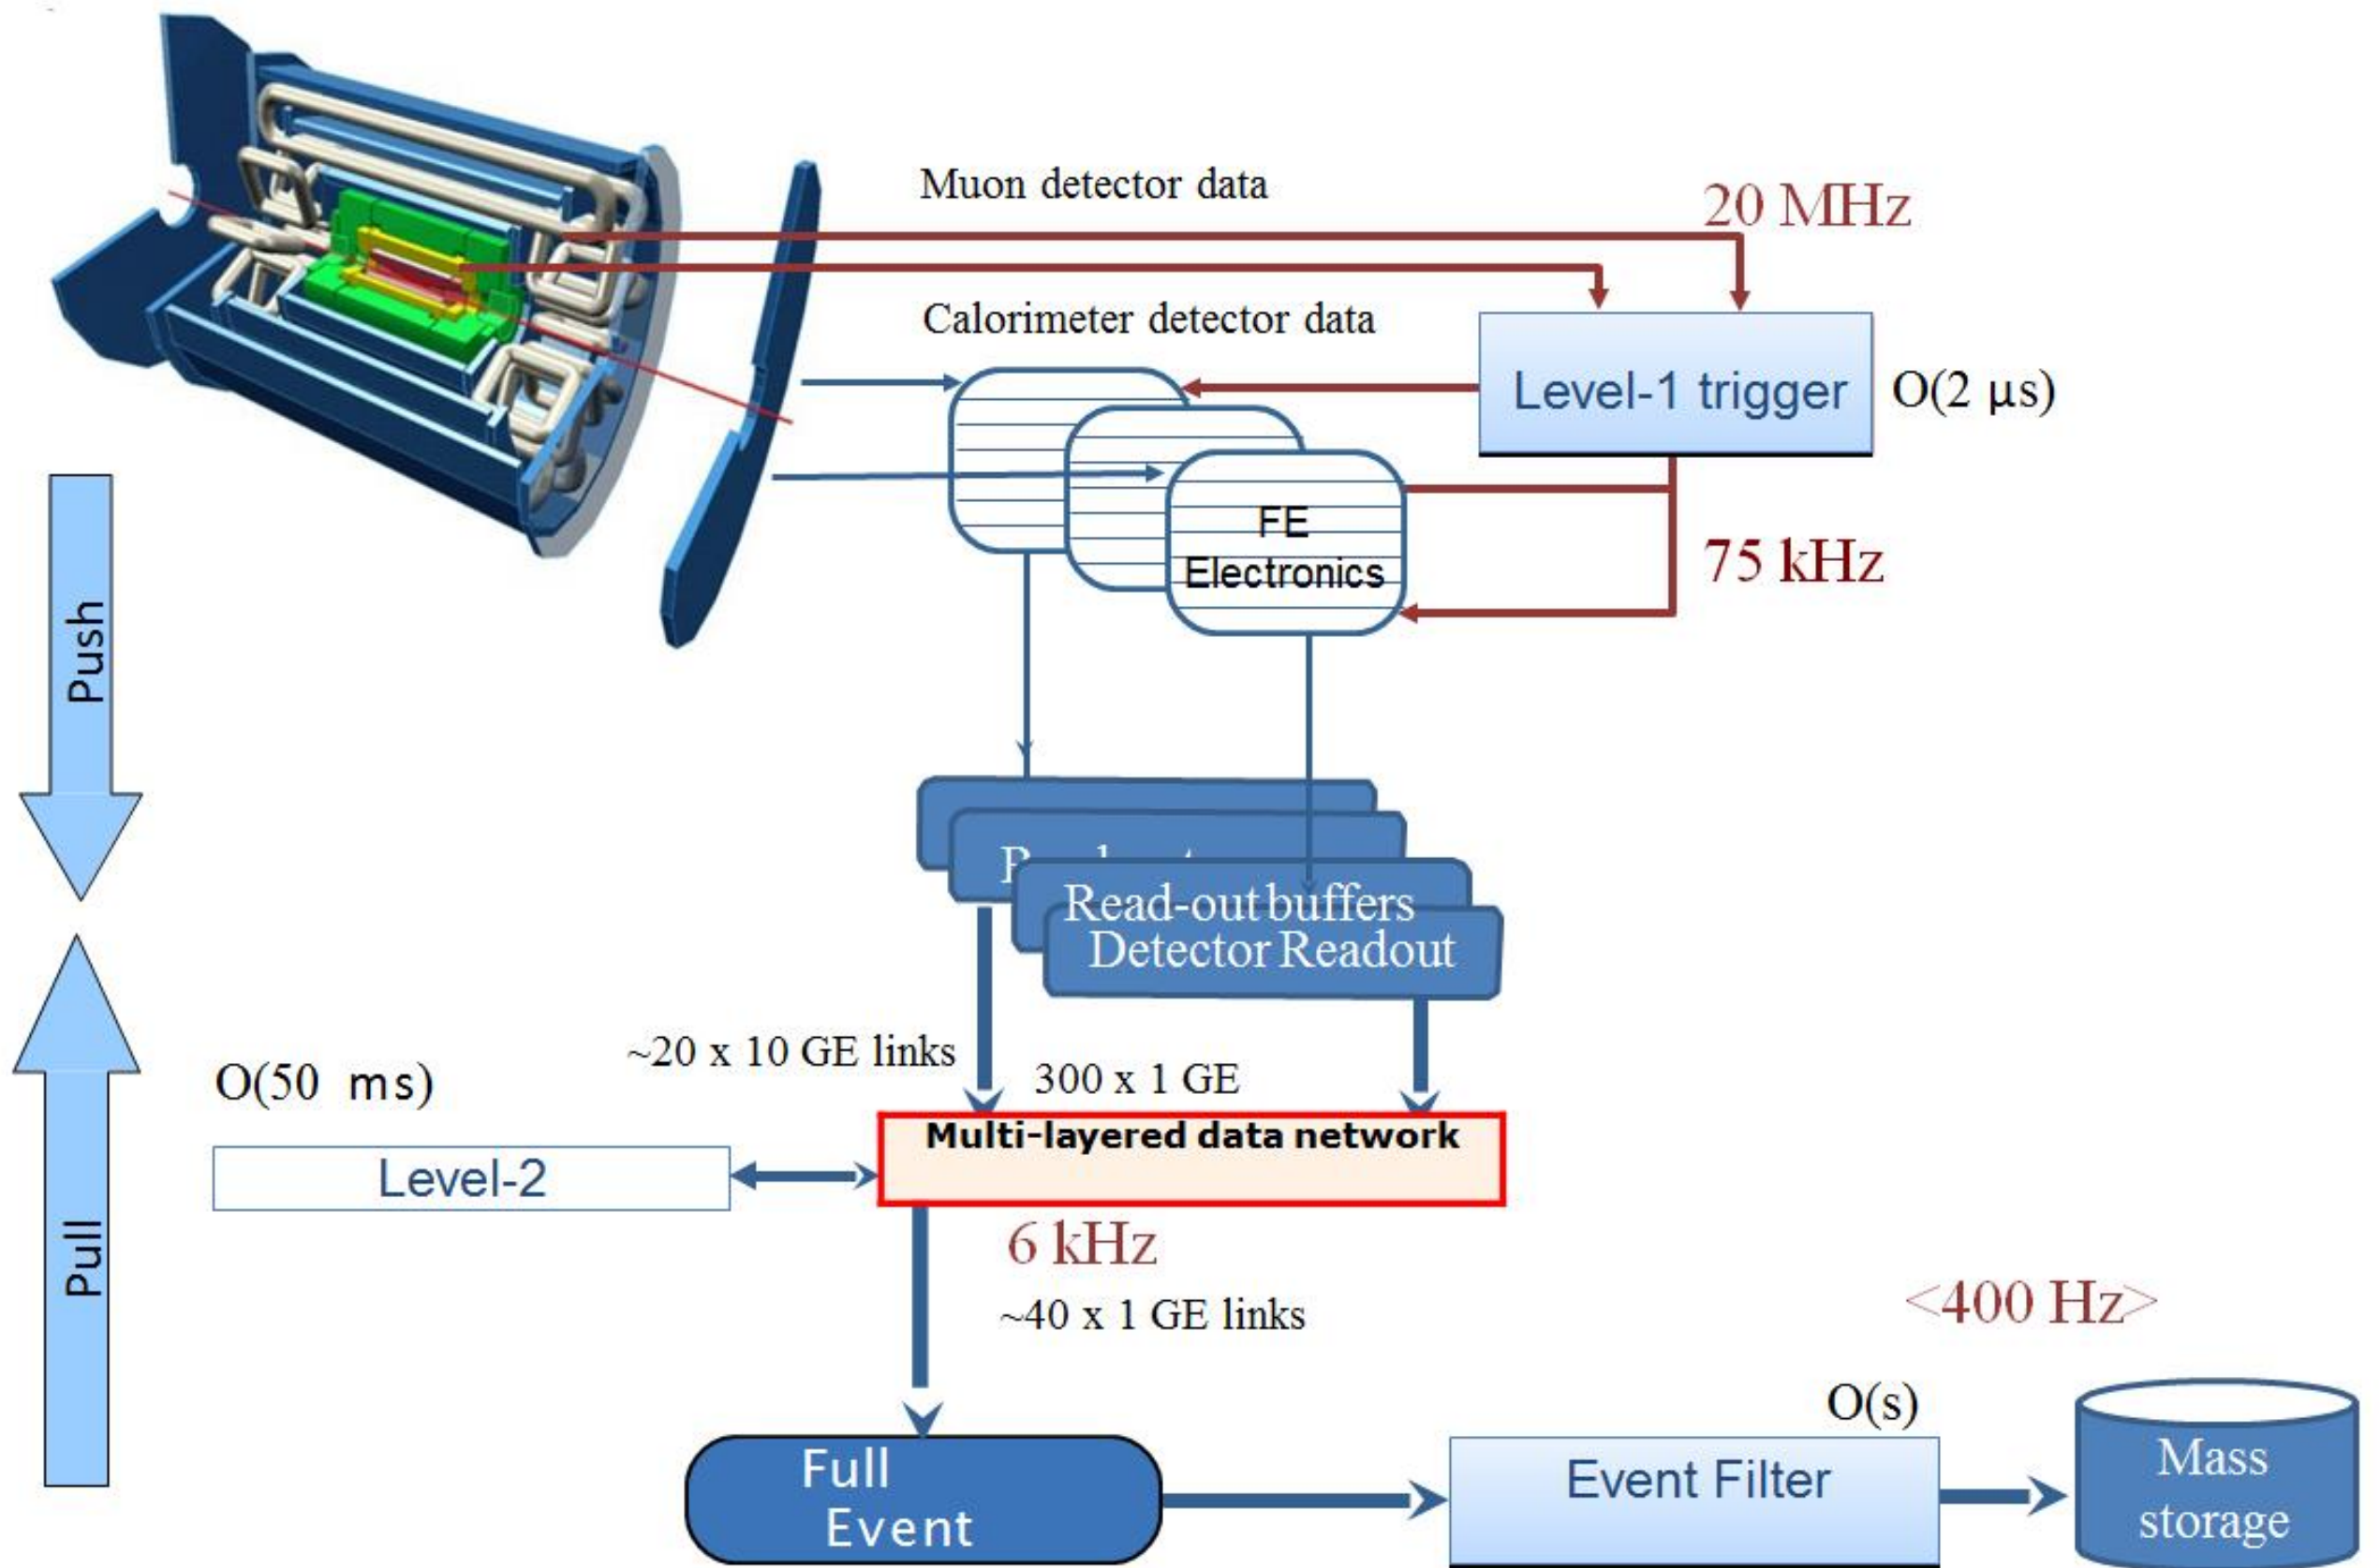
- Driving Parameters

	# of Trigger levels	Level-xRate (kHz)	Event Size (MB)	Network BW (GB/s)	Storage GB/s	kHz
Run 1	3	Lvl-1 75 HLT ~0.4	~1	10	0.5	~0.4
Run 2	2	Lvl-1 100 HLT 1	~2	50	1	1
Run 3	2	Lvl-1 100 HLT 1	~2	50	1	1
Run 4	3	Lvl-1 400 HLT 10	~5	2000	25	10

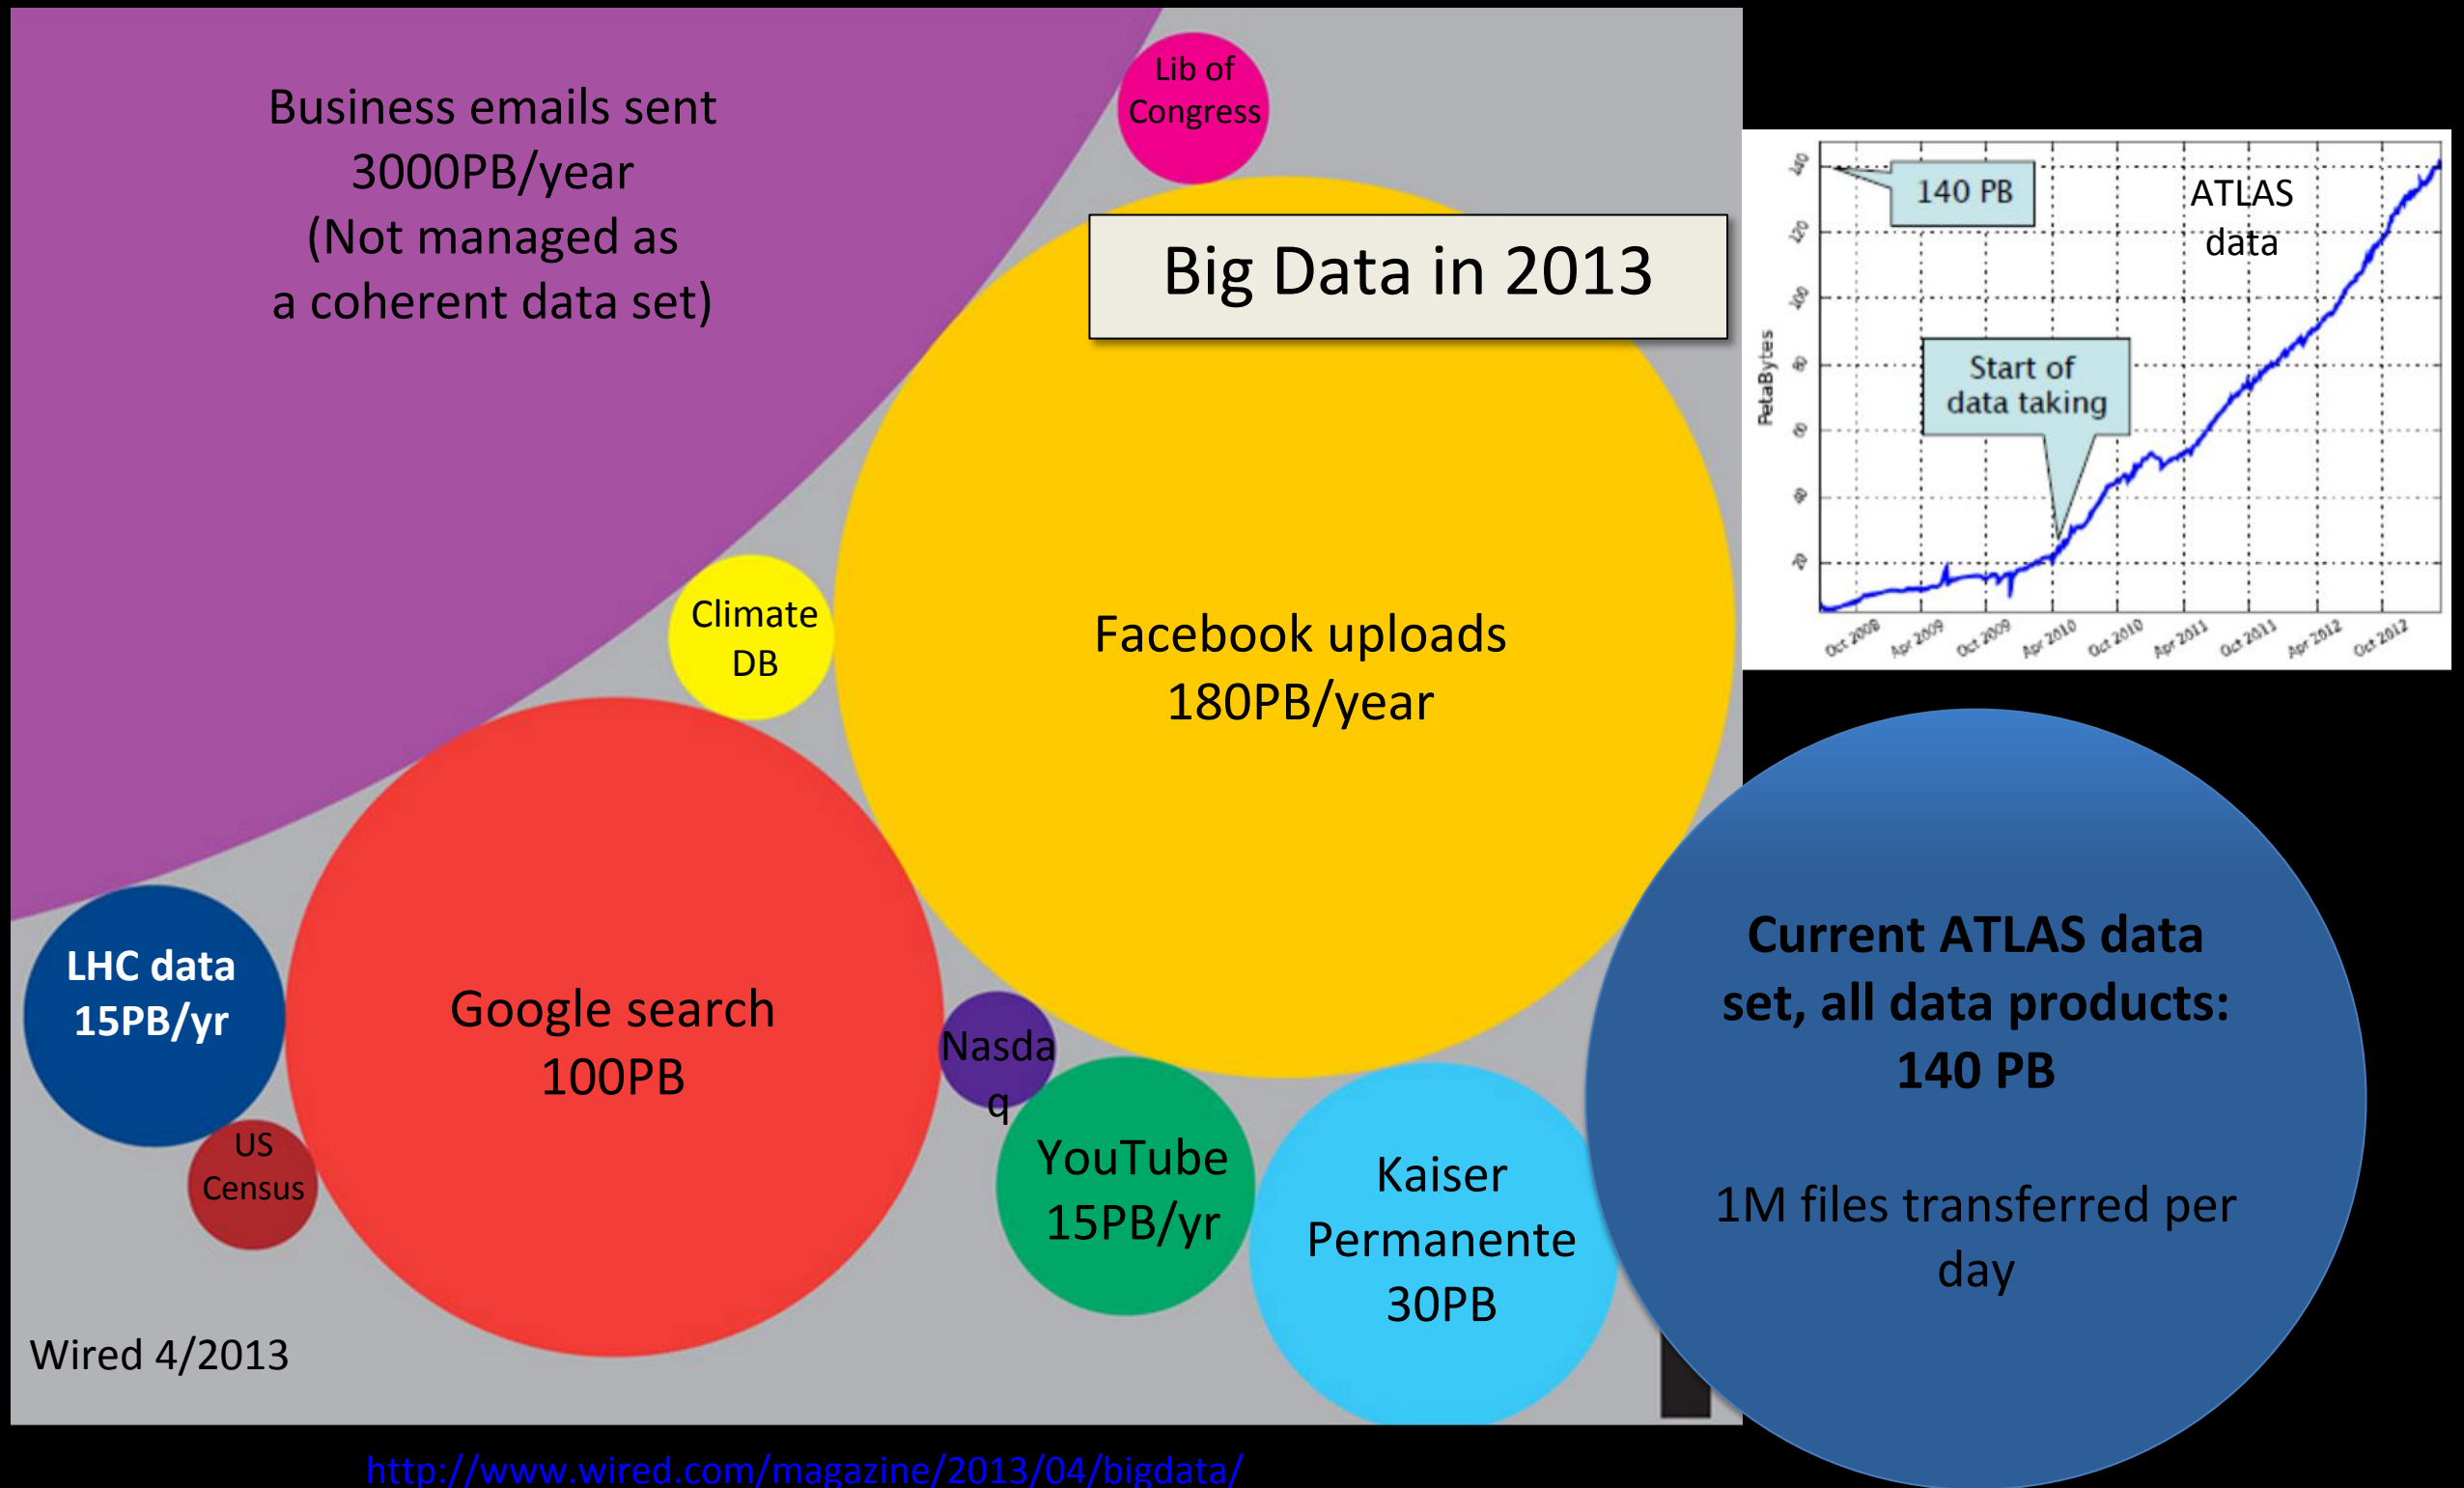
Notes

- Storage bandwidth includes compression factors
- Run 3: No major upgrades in the DAQ systems currently foreseen
- Run 4: Two stage hardware trigger (Level-0 and Level-1)

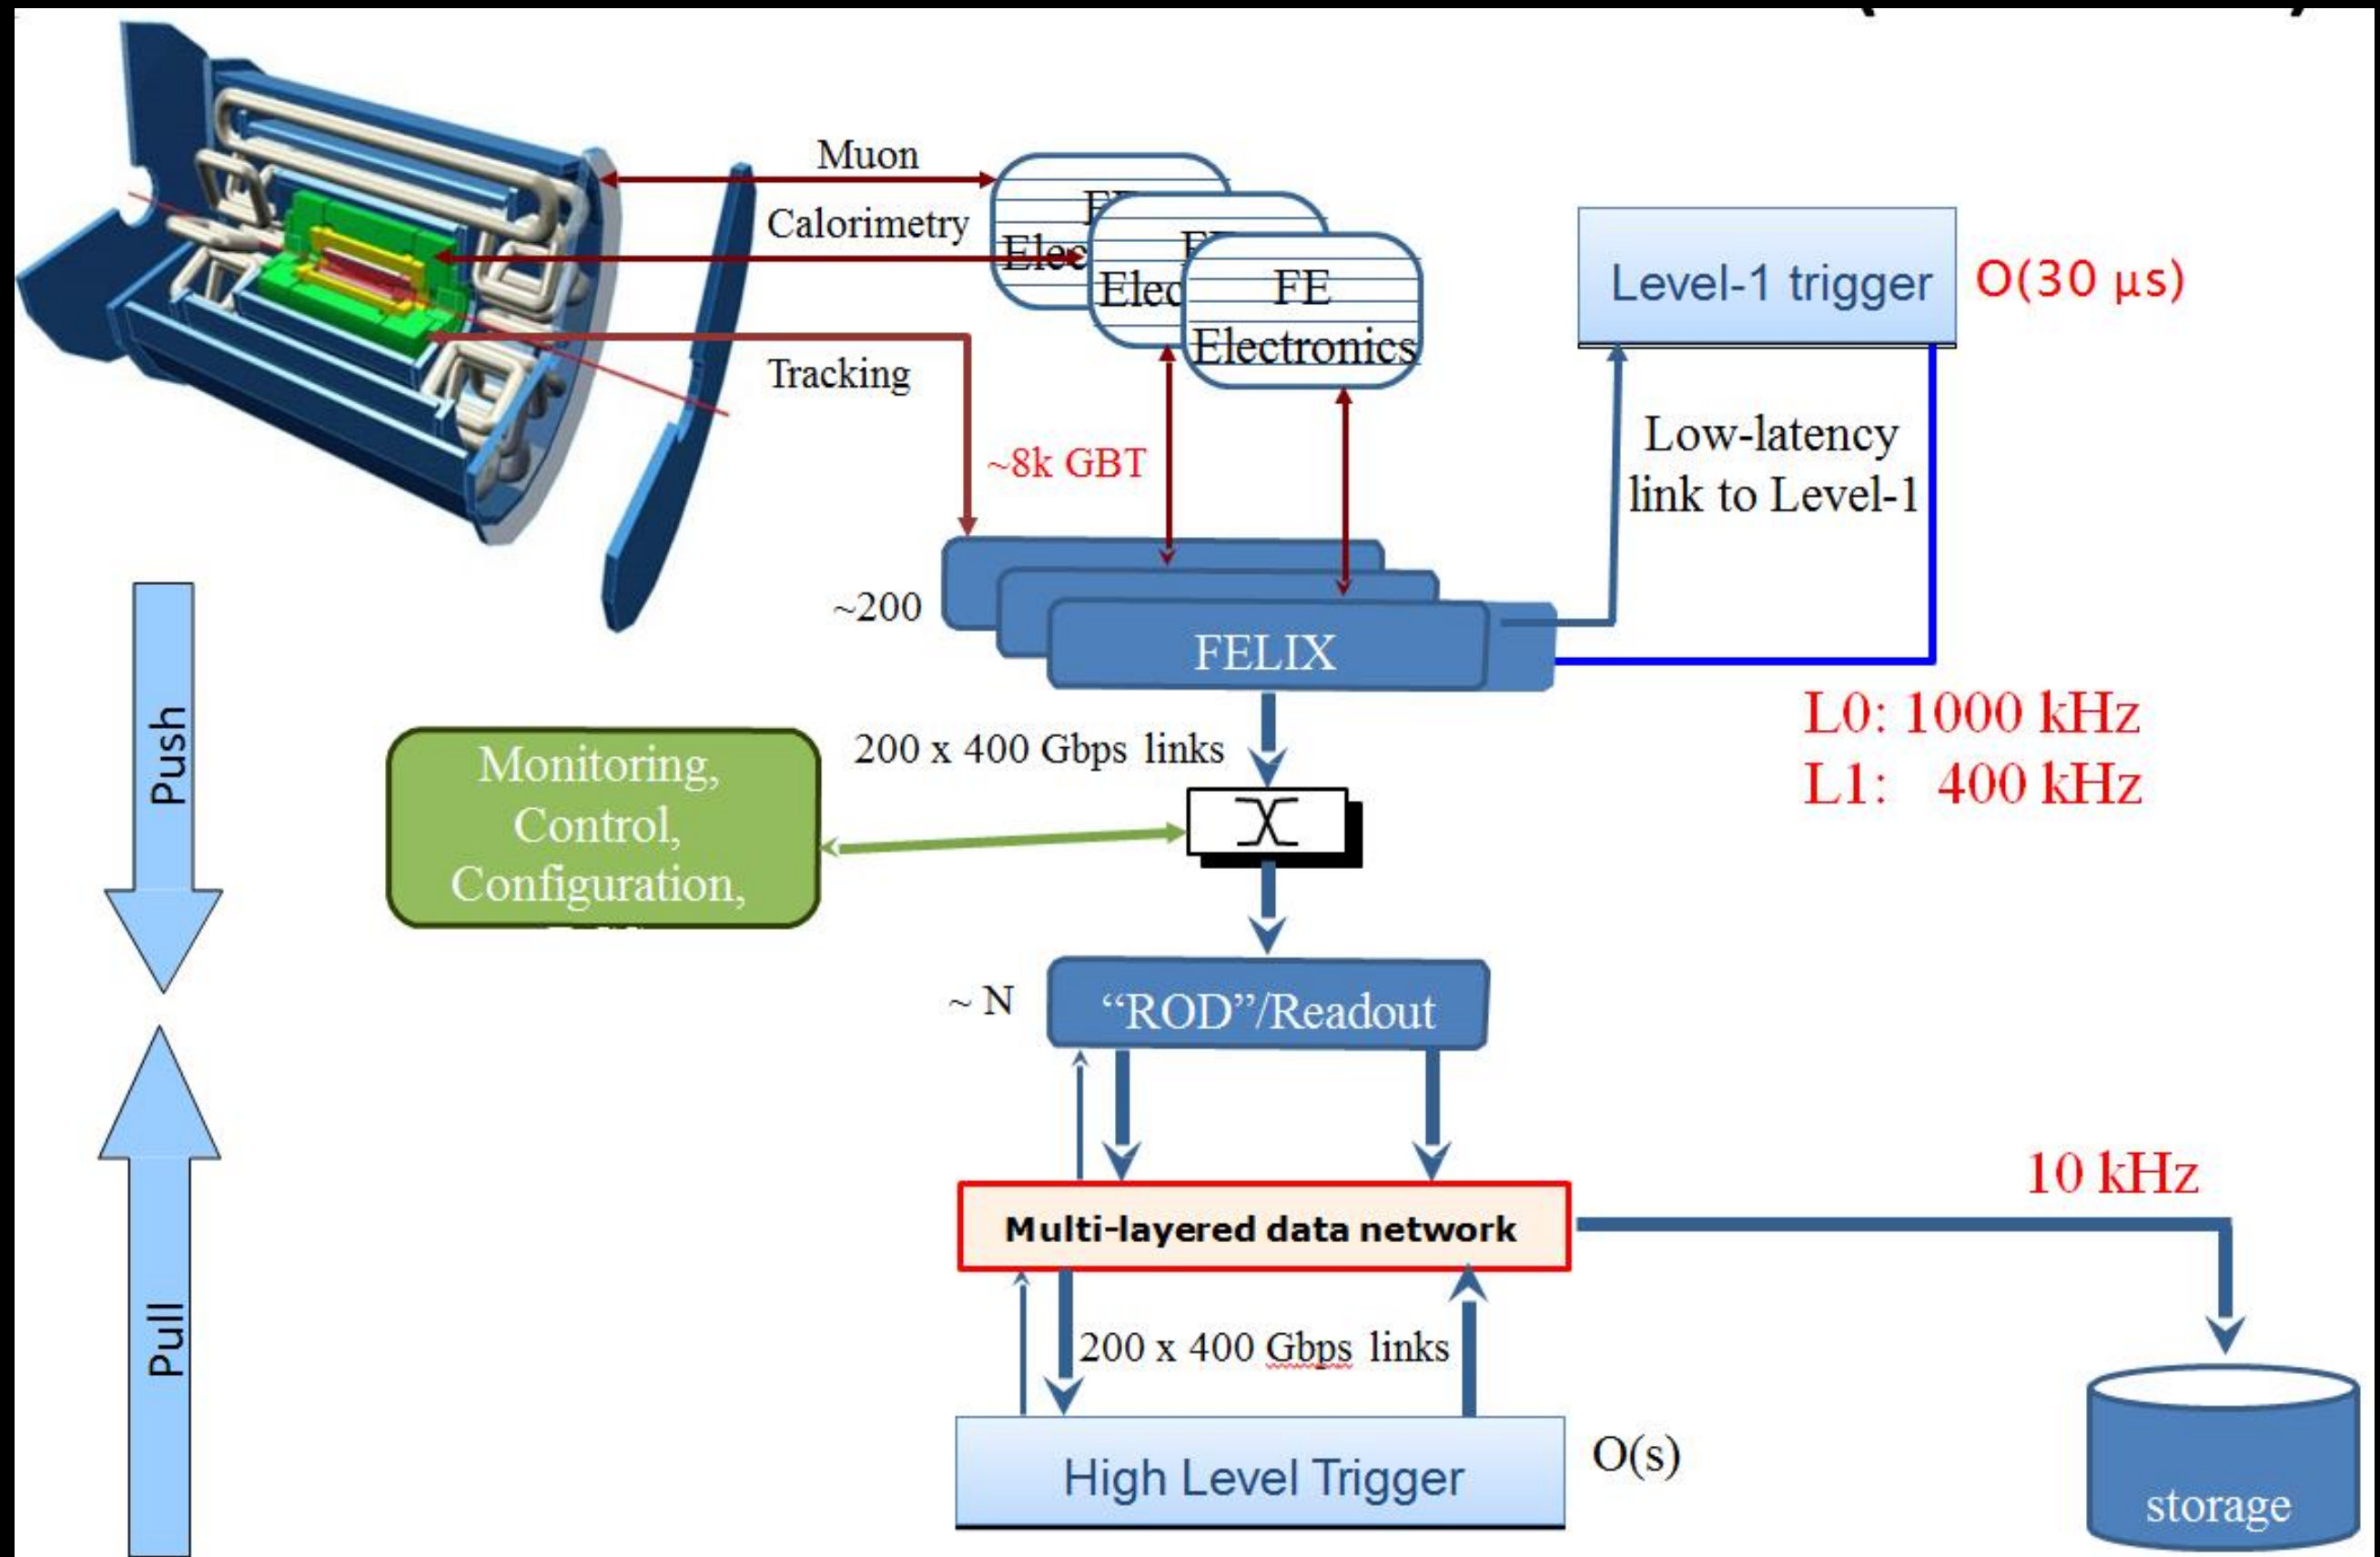
ATLAS Run 1



ATLAS is Big Data



Architecture in Run 4 (~2025)



Content Delivery Networks: Compare to Netflix

- HEP problem harder than Netflix?
 - Netflix delivers streaming video content to > 20M subscribers
 - Routinely quoted as the single largest user of bandwidth in the US
 - More than 30% of the traffic
- HEP has a different working point:
 - < # clients,
 - < distribution,
 - > bandwidth per client
- However, **much** larger data set
 - HEP can't make many multiple static copies
 - need different strategies instead:
 - make **dynamic** replicas and clean up when no longer useful
 - access data directly over the wide area networks


	NETFLIX	LHC Computing
Bandwidth per client	1.5Mbit/sec	1MByte/sec
Clients	1M*	200k cores
Serving	1.5Tbits	0.8Tbits
Total Data Distributed	12TB	140PB
Annual Budget	>\$4B	< \$.04B

Similar Problems:
Not all files
are equally
accessed



e.g. Forward
Physics ;-)

The Problem....




...the challenges
of computing at
the *Petabyte*
scale

*How to make a torrent
of data available for
science in the most
efficient, transparent,
and cost-effective way?*



1980s: Plethora of architectures & OSES




1990s: **Uniform OS/architecture**
Linux/x86 standard for commodity cluster computing




2000s: **Uniform fabric and access**
Globally federated resources enabled by **network and grid**

HEP computing evolution – growing uniformity
counters growing scale and complexity

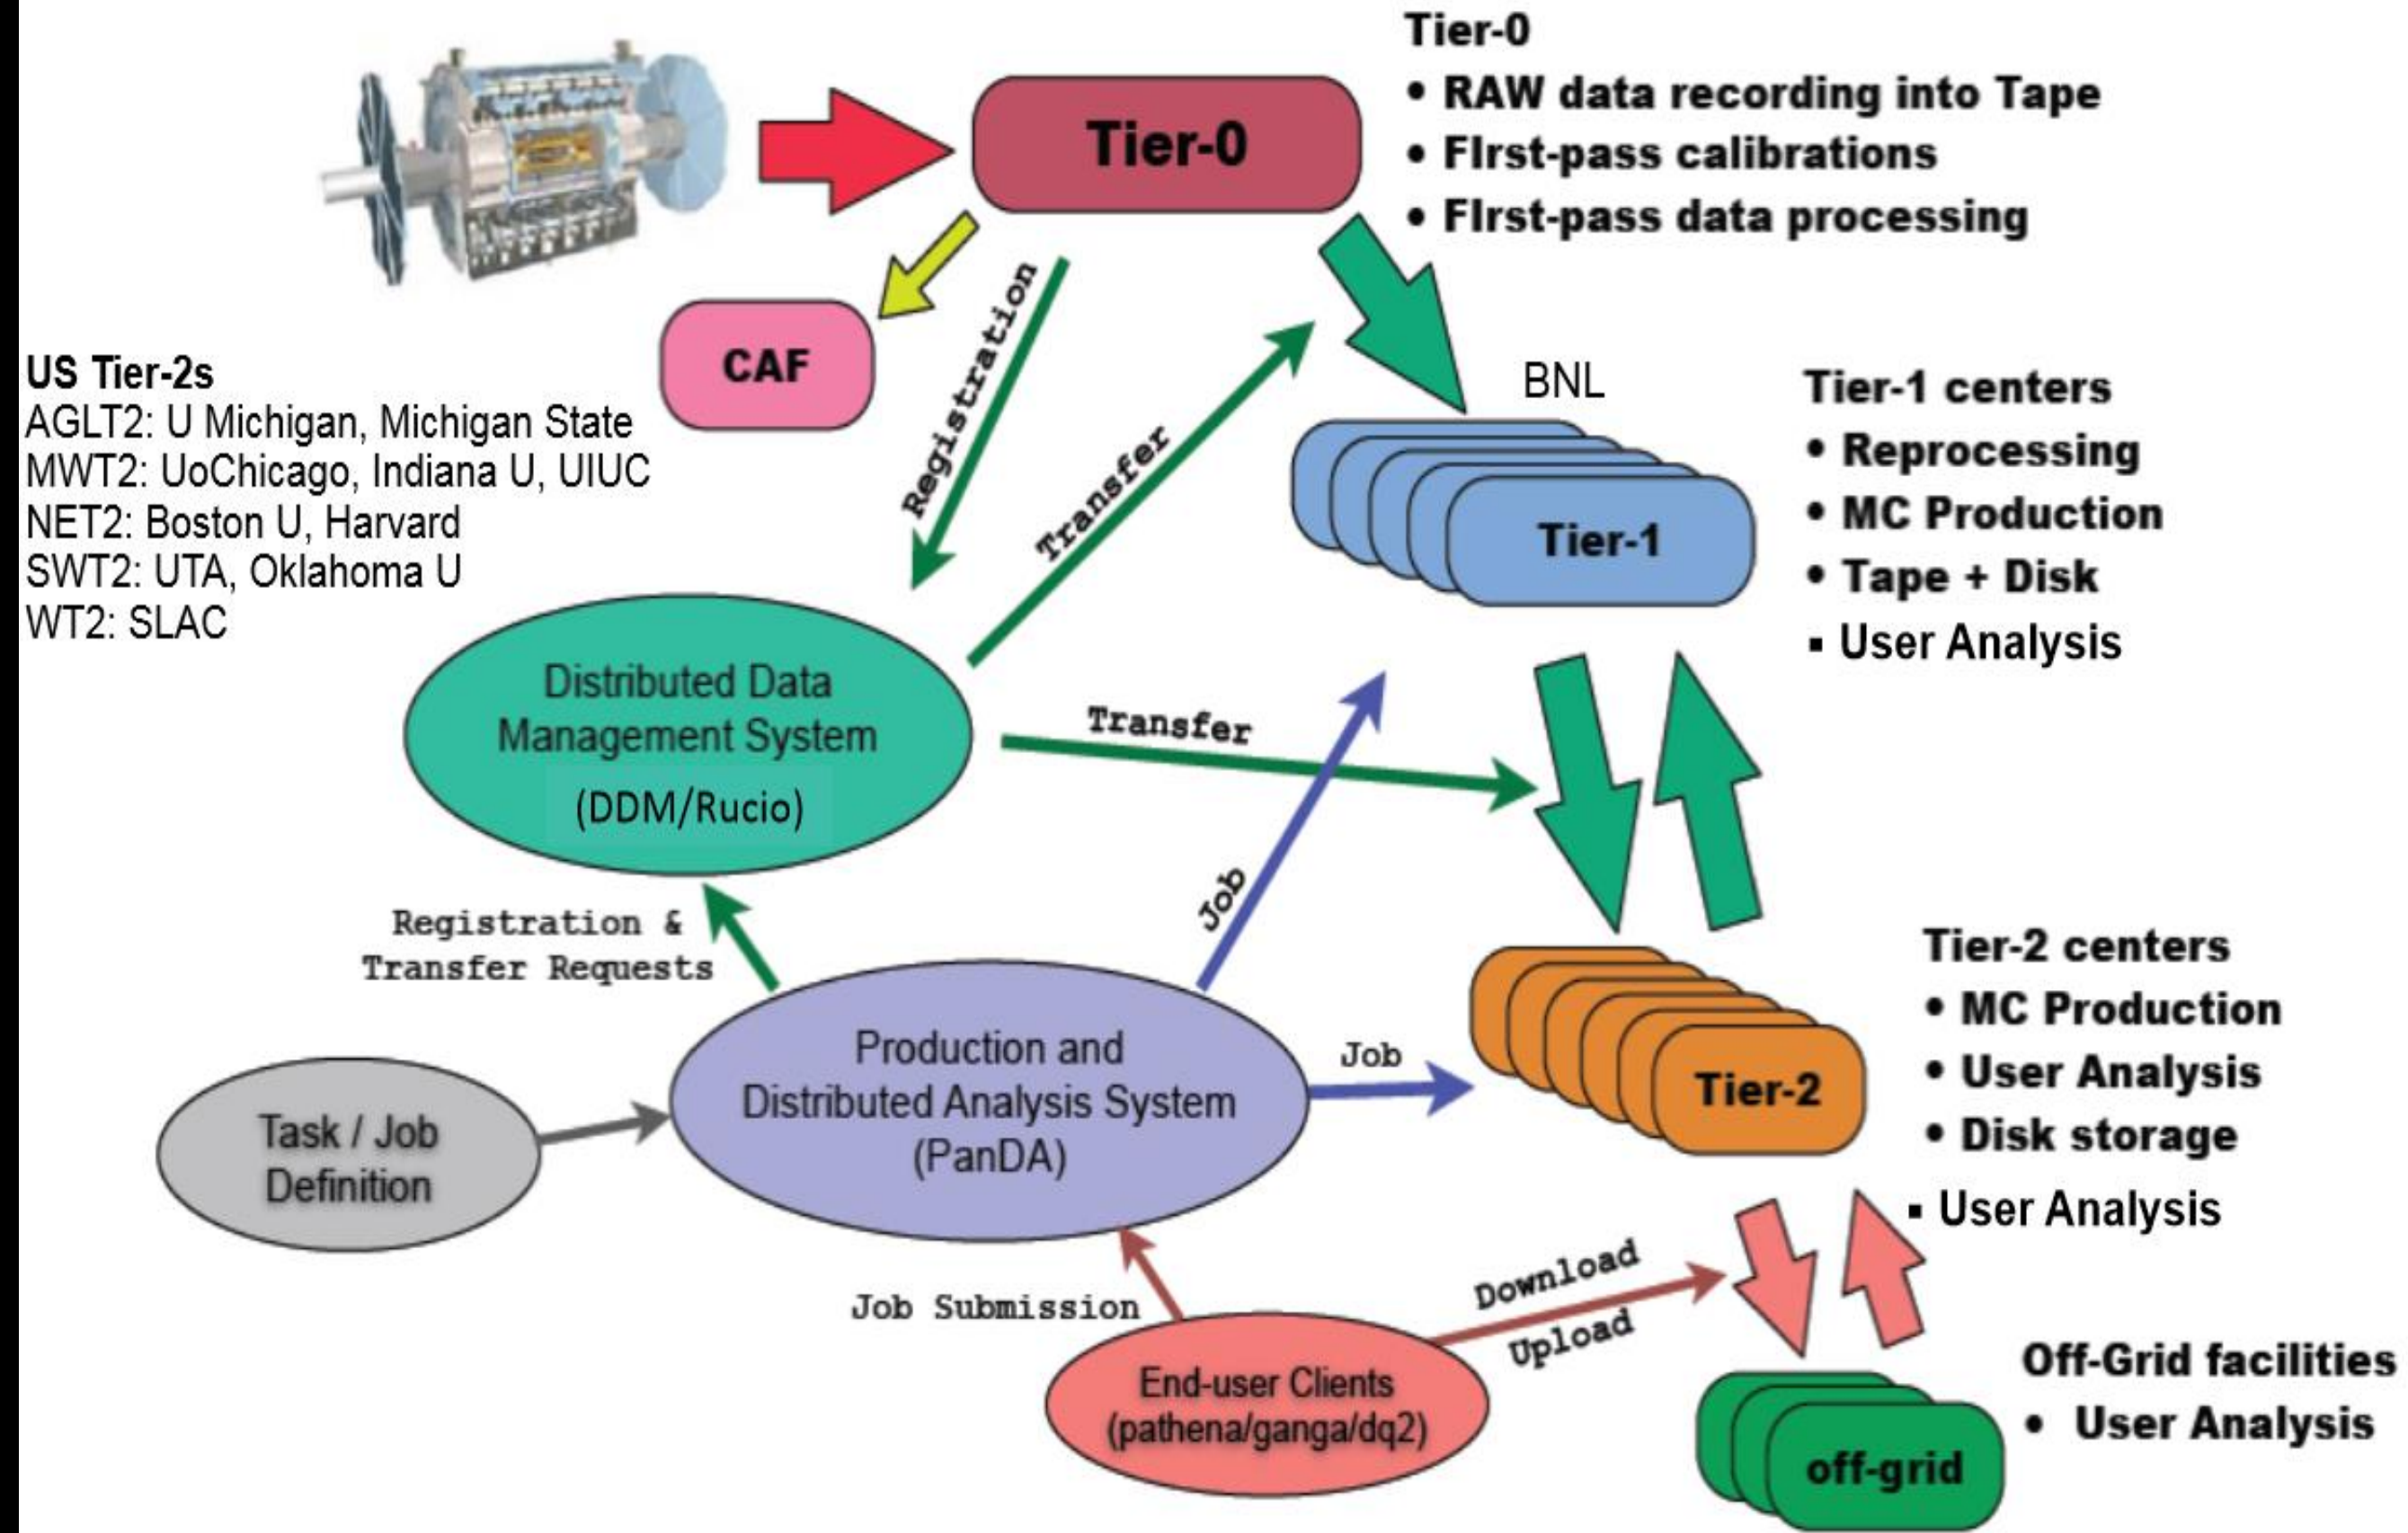


2010s: **Uniform environment**
VMs and clouds put the user in control of the environment – take it with you anywhere and everywhere



2010s: **Uniform data access**
Working towards transparent distributed data access enabled by the **network and web**

ATLAS Distributed Computing ...




... Using a Global Computing Infrastructure

Tier-0 (CERN): data recording, reconstruction and distribution

Tier-1: permanent storage, re-processing, analysis

Tier-2: Simulation, end-user analysis



Nearly 160 sites in 35 countries

~350,000 cores

200 PB of storage

> 2 million jobs/day

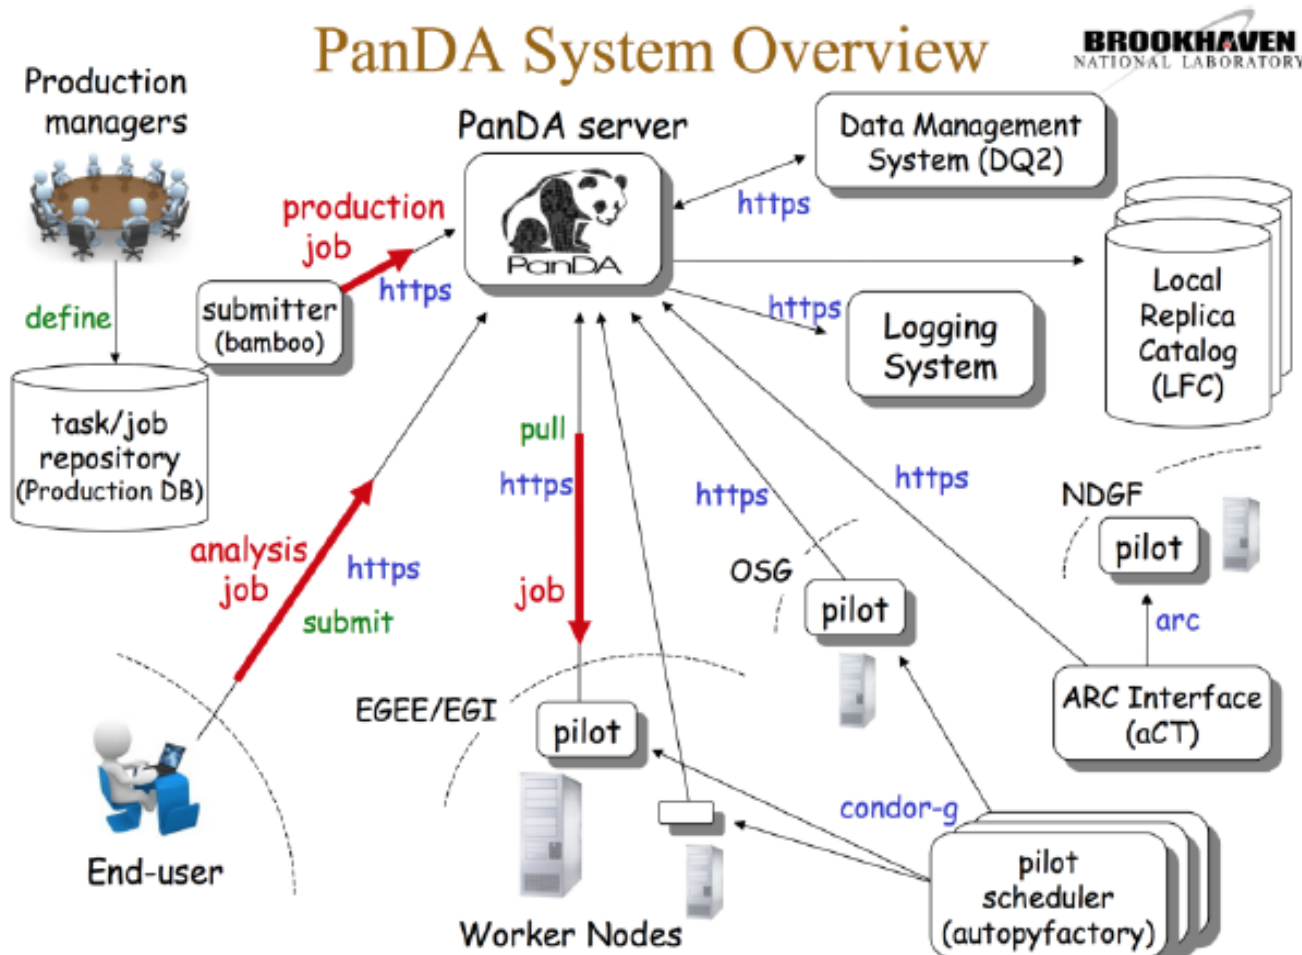
10 and 100 Gb links

The Worldwide LHC Computing Grid

WLCG: An international collaboration to distribute and analyze LHC data. Integrates computer centers worldwide that provide computing and storage resources into a single infrastructure accessible by all LHC physicists

PanDA Workload Management System

PanDA System Overview



BNL's Physics Applications Software (PAS) group leads development of the PanDA workload management system with UT Arlington
PanDA manages processing and data workflows for large scale data intensive computing

- 2005: Initiated for US ATLAS
- 2008: Adopted ATLAS-wide
- 2009: First use beyond ATLAS
- 2012: ASCR/HEP funding for Exascale BigPanDA
- 2013-14: New PanDA based ATLAS prod system
- 2014: New Event Service fine grained processing
- 2014: PanDA community growing in HEP, NP, cosmology... US and international


Global ATLAS PanDA operations:

Up to ~200k concurrent jobs, 1.5M/day, O(100M)/year
~1400 ATLAS users, ~140 sites

1.2 Exabytes per year processed (2013)


PAS also does software infrastructure, cloud & HPC porting, offline software development, leveraging new processor technologies, software project management... for BNL's HEP programs in all Frontiers, and the wider community

~ 180 k running jobs at peak



RACF at BNL: An Overview


- Formed in the mid-1990's to provide centralized computing resources for the four RHIC experiments (BRAHMS, PHOBOS, STAR, PHENIX)
- Role was expanded in the late 1990's to act as the US Tier-1 computing center for the ATLAS experiment at the LHC
- Small but growing neutrino and astrophysics presence (Daya Bay, LBNE, LSST)
- Located in the Brookhaven Computing Facility
- 30 FTEs providing a full range of scientific computing services for more than 4000 users



RACF: Setting The Scale

□ RHIC


- **1300** COTS Compute Servers (**32k** logical CPU cores)
- **16 PB** of Distributed Disk Storage on Compute Nodes, >1500 Gb/s aggregate data transfer rates observed between servers
- **5 Robotic Tape Libraries** w/ 50+ tape drives and 50k cartridge slots



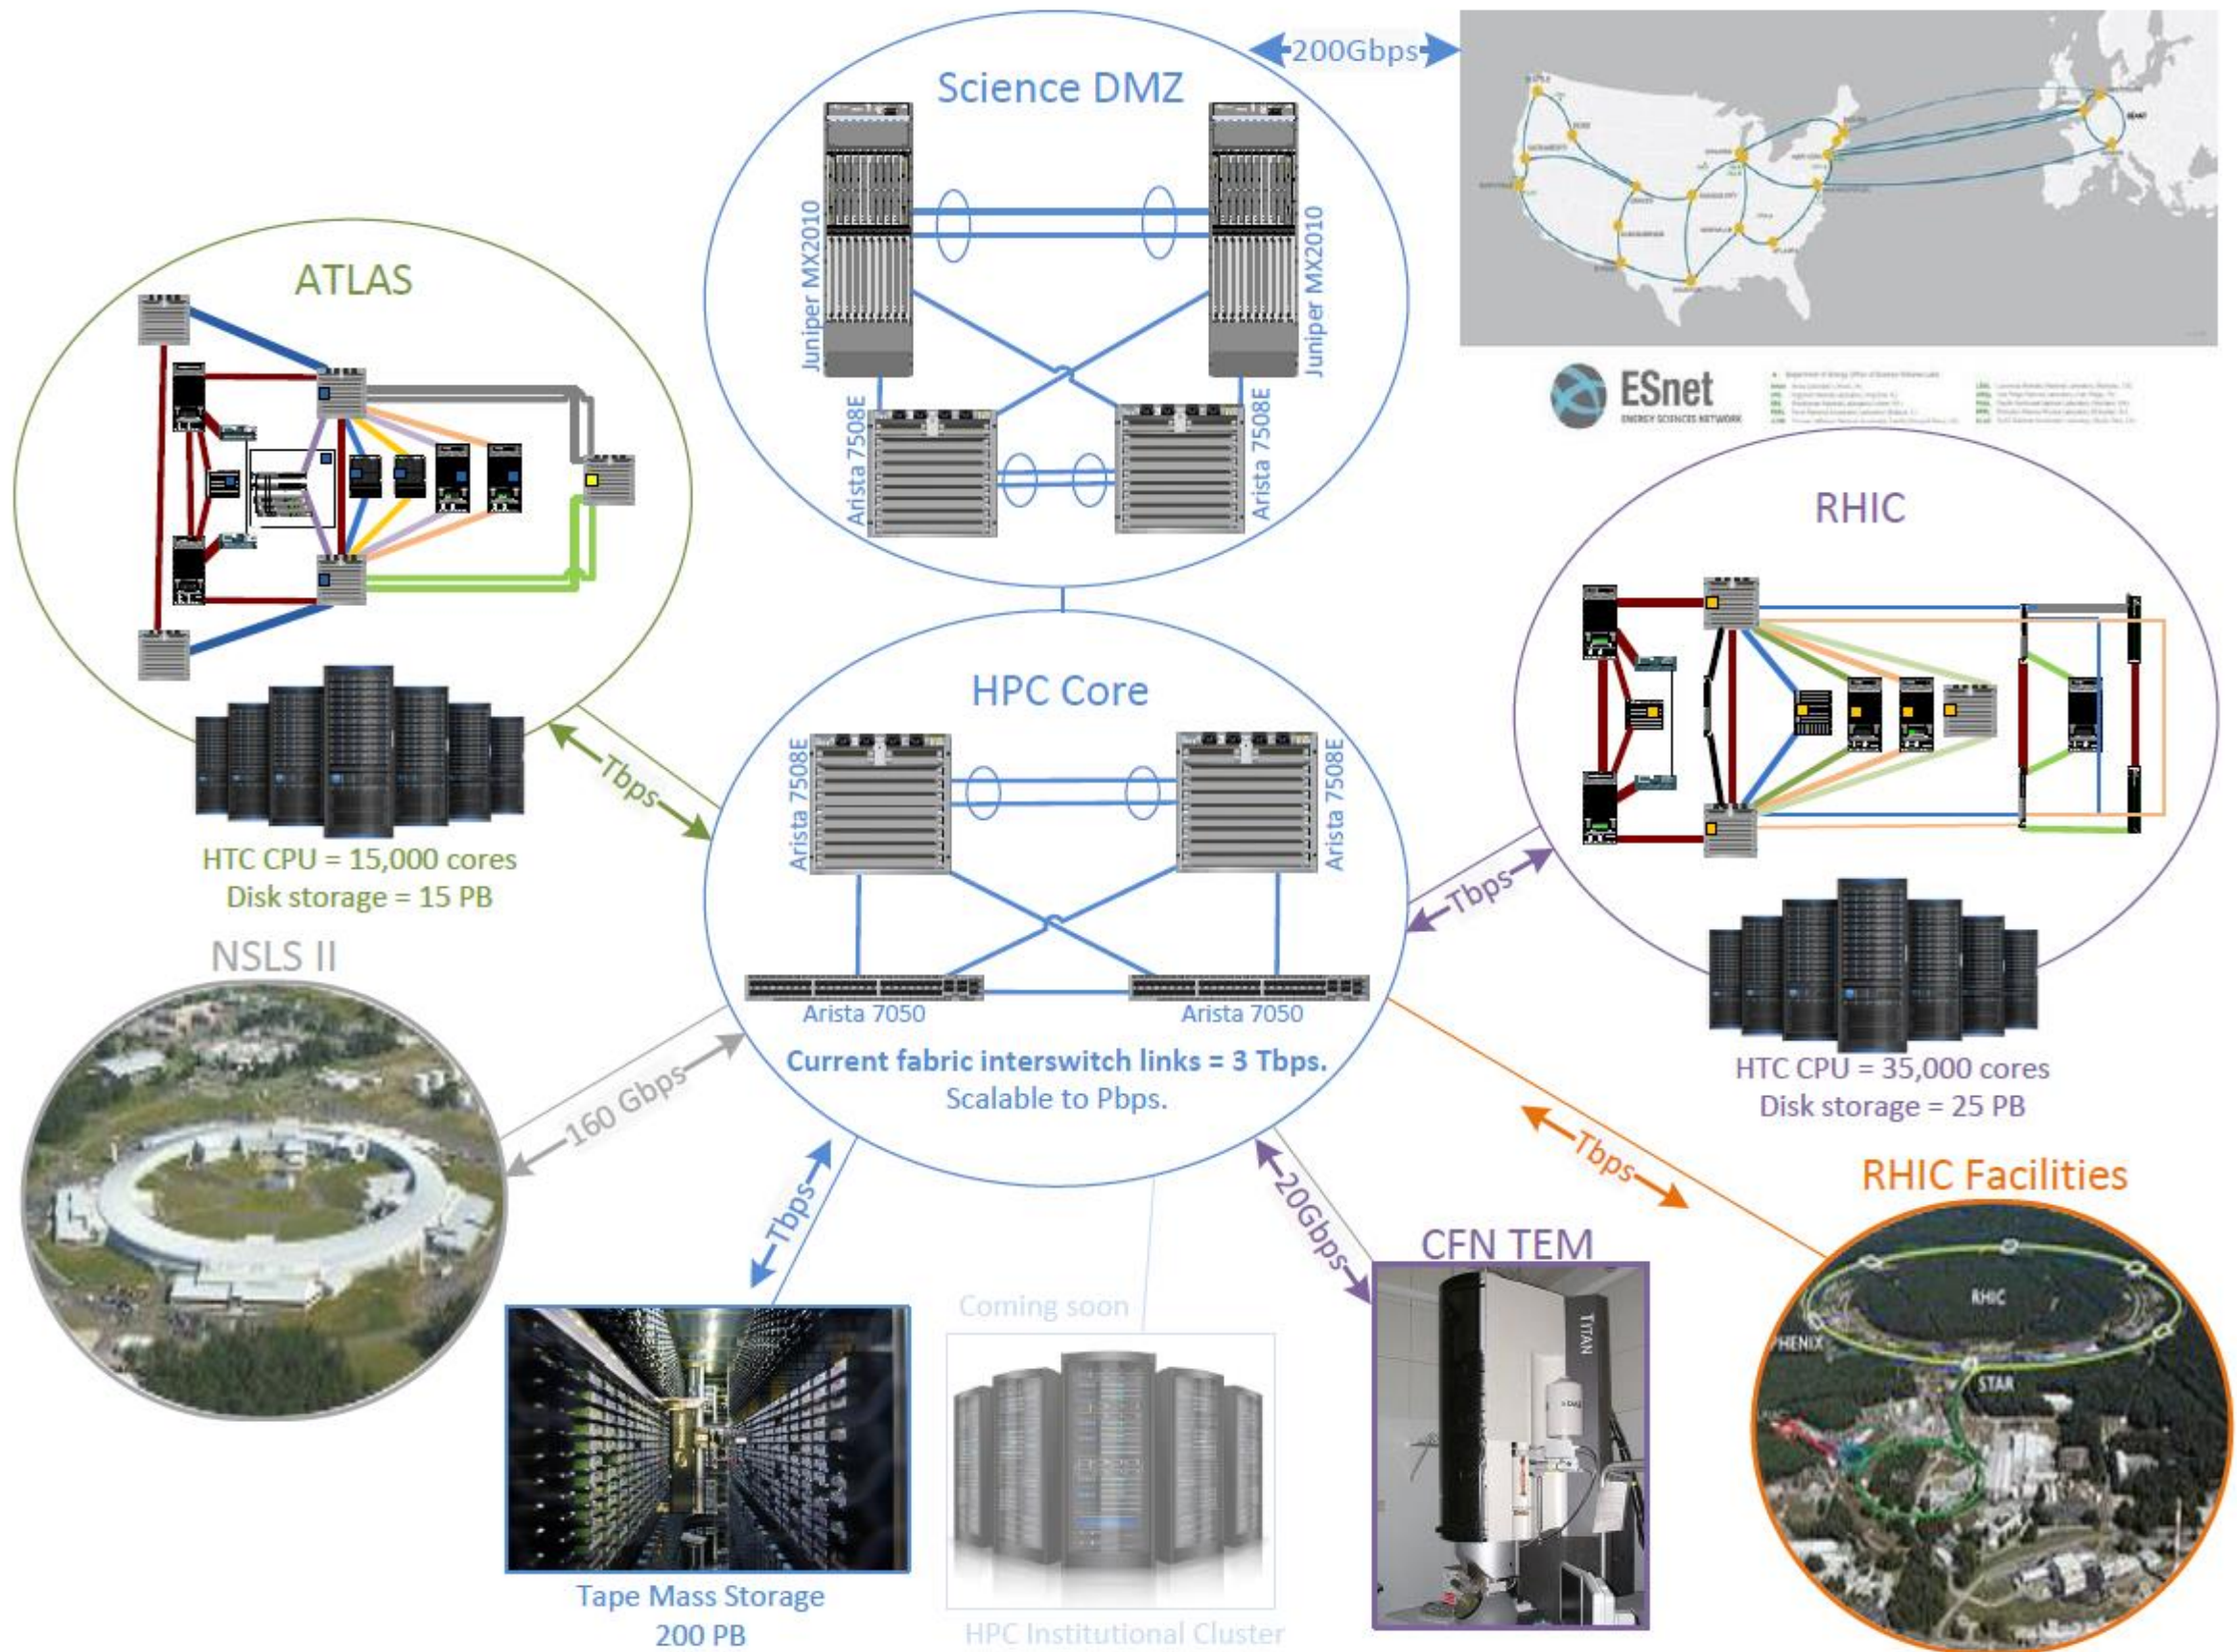
RACF: Setting The Scale

□ ATLAS

- 830 COTS Compute Servers (17,000 CPU cores)
 - 42 Dedicated Servers providing 13.6 PB of Disk Storage
 - 4 Robotic Tape Libraries w/ 40 tape drives and 36k cartridge slots
- ## □ Magnetic Tape Archive (200 PB Capacity)
- Data inventory of 60 PB on 58,000 tapes



RACF – Setting the Scale



Our Infrastructure is Services

A rich environment of common services that can be flexibly composed to meet specific requirements of science domains across BNL and DOE SC

Human Computer Interaction

Platform
Instantiation Interface

Workflow
Composition &
Execution Manager

Visual Analytics
Interface

Analytic Services

Data Mining
Semantic Analysis
Data Fusion

Data Transfer
Tools

Data Services

Metadata Harvesting &
Management

Indexing,
Discovery &
Dissemination

Simulation Services

Simulation
Frameworks

Scalable Debuggers
Scientific Libraries

System Software & Middleware Services

MPI
ADIOS

Map
Reduce

HIVE

Key Value
Stores

Graph
Databases

SQL
Databases

Message
Queues

SDN

Infrastructure Services

HPC
Compute

Utility
Compute

Advanced
Networking

Parallel File
Systems

Network
Storage

Archival
Storage

Object
Storage

Visualization
Environments

Workflow Composition

Security




RACF Area of Expertise



ATLAS Workload – Managed by PanDA

dashboard



Global ATLAS PanDA operations


Up to ~200k concurrent jobs
25-30M jobs/month at >100 sites
~1400 ATLAS users, ~140 sites

First exascale workload manager in HEP
1.2 Exabytes processed in 2013

1.1 Exabytes in 2014


Exascale scientific data processing today

ATLAS Processing Contribution by Tier-1 Site in 2015



Provisioning for Peak Demands


- The “dream” of short turn-around times for workflows
 - ✦ Short latencies in particular in analysis workflows are important for science efficiency
 - ✦ Use resources from a larger pool when they are needed, should also result in more cost-effective solutions
- Separating the processing and storage services allows them to scale independently
- e.g. ATLAS and CMS are looking at ways to double available resources for periods of time
 - ✦ using Amazon services



**Provisioning for peak requires
that we use pooled resource
—> Clouds or large HPC Center!**

Provisioning for Peak Demands

- The “dream” of short turn-around times for workflows
 - ✦ Short latencies in particular in analysis workflows are important for science efficiency
 - ✦ Use resources from a larger pool when they are needed, should also result in more cost-effective solutions
- Separating the processing and storage services allows them to scale independently
- e.g. ATLAS and CMS are looking at ways to double available resources for periods of time
 - ✦ using Amazon services


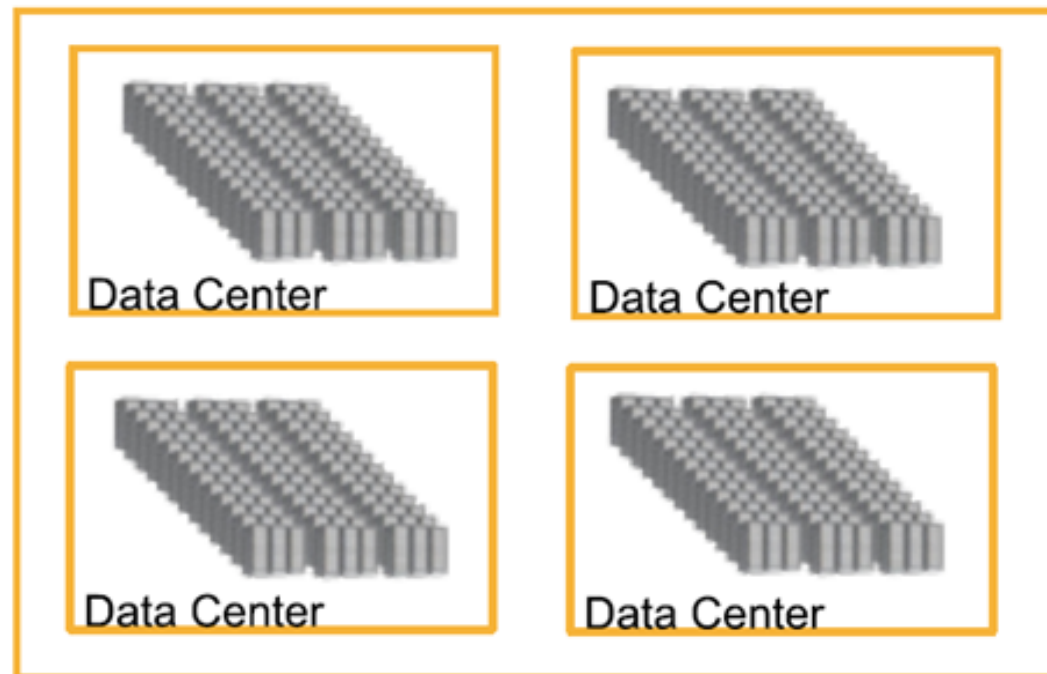


**Provisioning for peak requires that we use pooled resource
—> Clouds or large HPC Center!**


The Potential of Commercial Cloud Resources

A Snapshot of AWS Global Capacity

J. Kinney (AWS)



- 53 AWS Edge Locations
- 11 Regions
- 28 Availability Zones
- 2 or more AZs per Region
- 1-6 Data Centers per AZ
- 50,000-80,000+ servers per DC
- Up to 102 Tbps provisioned to each DC



Cost for Compute (AWS Spot) quickly approaching cost for dedicated resources

- A cost-efficient way to serve peak demand

Running ATLAS Jobs in the Cloud – At Scale and at Low Cost

Joint Project (Amazon AWS, BNL/ATLAS, ESnet) to investigate feasibility (technically and financially) of large-scale usage of commercial cloud resources

AWS: Provide expertise and guidance to BNL/ATLAS, credits for AWS service investigation and scale-out tests

BNL: Provide ATLAS-compatible VM image and provisioning infrastructure, incl. demand-driven (i.e. via PanDA server API) VM lifecycle management (create, retire, terminate)

ESnet: Provide high-performance (up to 100G) network connectivity between AWS facilities and sites connected to R&E networks (general peering and AWS Direct Connect)

Has made AWS partially waive Egress traffic fee (at level of 15% of total bill)

Cost of AWS/EC2 spot slightly lower than dedicated farm resources at BNL

Running ATLAS Jobs in the Cloud – Experience

“Unlimited” spot resources available on demand

Had no issues ramping up quickly from 1000 to ~6000 8-core instances (several instance types) in a single (out of 3) AWS region in the U.S.

Ran at level of 6000 instances for a few days with very low fluctuation (VM instance termination) due to spot overbidding

<1% of total running VM instances were terminated by AWS while production jobs were running during a multi-day period

Most of the terminations occurred within the first hour after VM instance creation -> no cost to us

Public cloud dvantages: “Unlimited” horizontal scaling in AWS EC2/S3 in terms of network bandwidth between compute and storage. Very high performance Object Store at low cost (when used as temporary storage for intermediate data products)

Running ATLAS Jobs in the Cloud – Potential going forward

Combination of “unlimited” capacity whenever we need it and competitive pricing makes AWS (and presumably other commercial cloud providers) an ideal resource to cover peak demands

Could think of deploying only components (kind and quantities) & services at our dedicated data centers where cloud providers cannot cope (yet, i.e. technical capabilities and cost)

Potential to vastly reduce size and scope of our dedicated (and aging) hardware deployment

Potential to lower computing facility operations cost at improved performance (whenever the collaboration is in desperate need) and availability (e.g. the availability of AWS services is much higher than what WLCG sites provide)

Potential to vastly increase our flexibility

Using cloud computing makes us nimble whenever we need specific resources/platforms – temporarily or for long periods.

Running ATLAS Jobs in the Cloud – Matching Workloads

But all these wonderful things are not compatible with and/or applicable to our current processing model

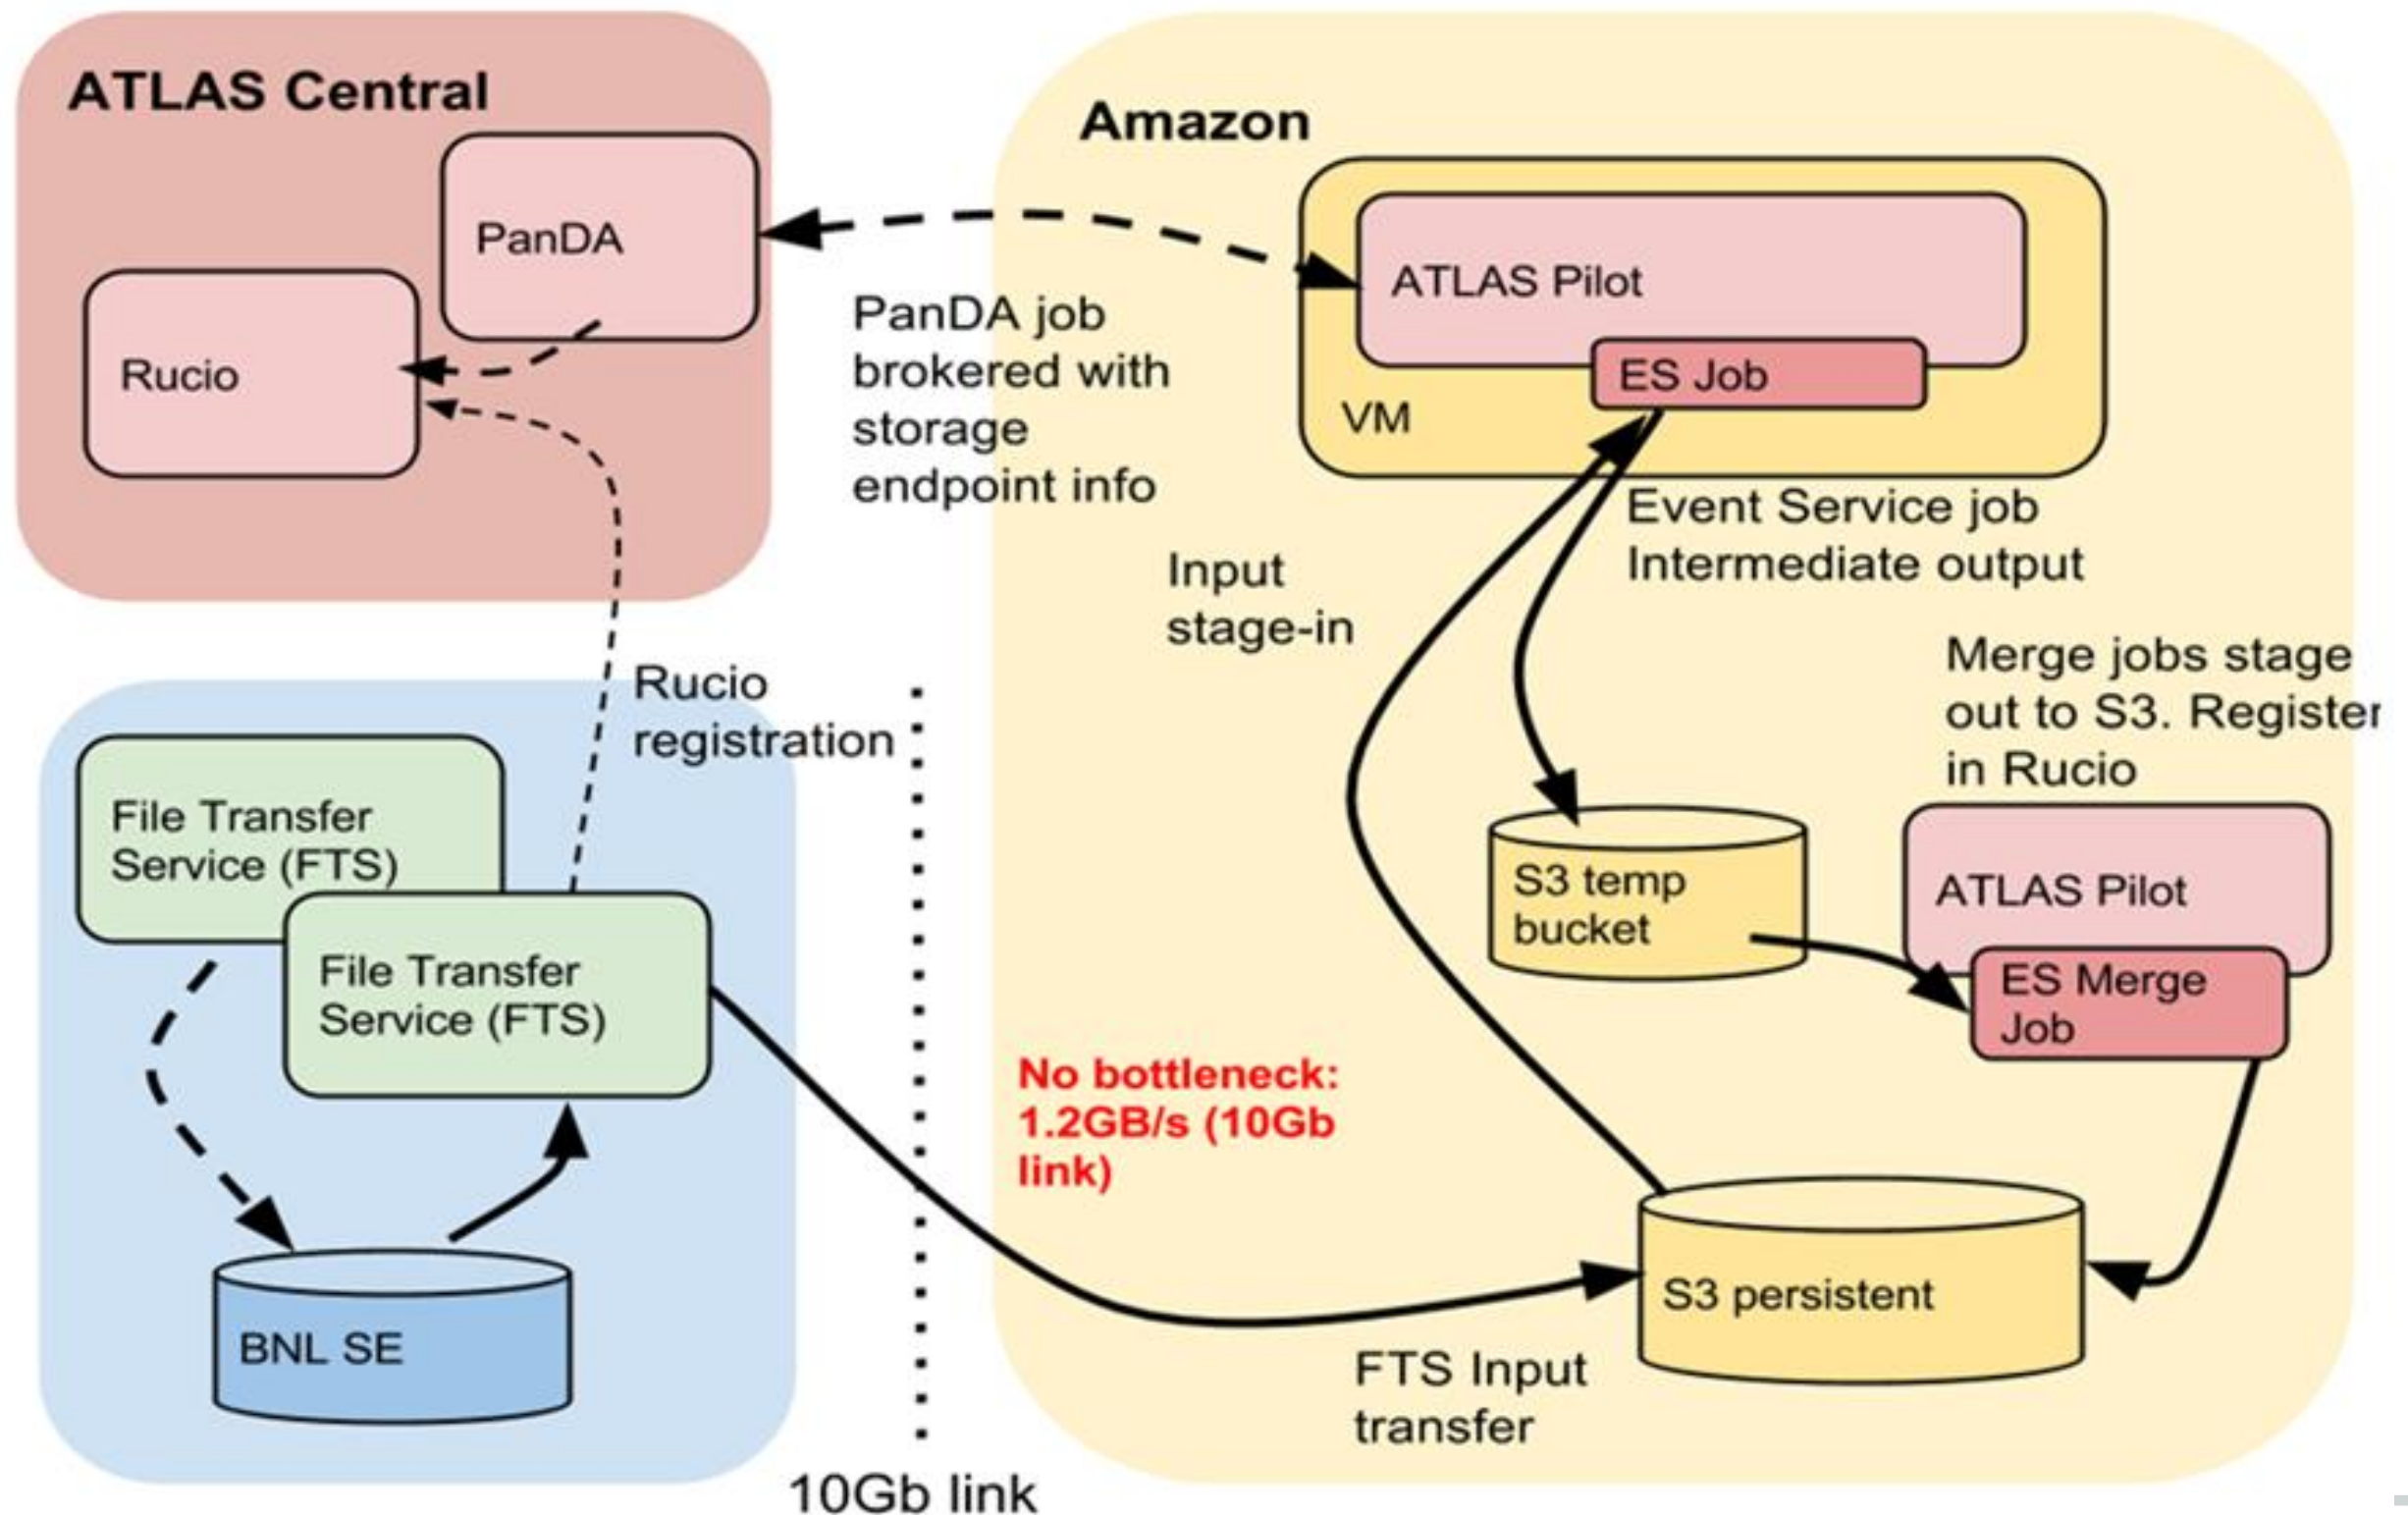
Most of our compute-intensive jobs run for 6-24 hours making spot VM instance terminations likely at probability of up to 80%

Potentially a huge waste of resources we would have to pay for

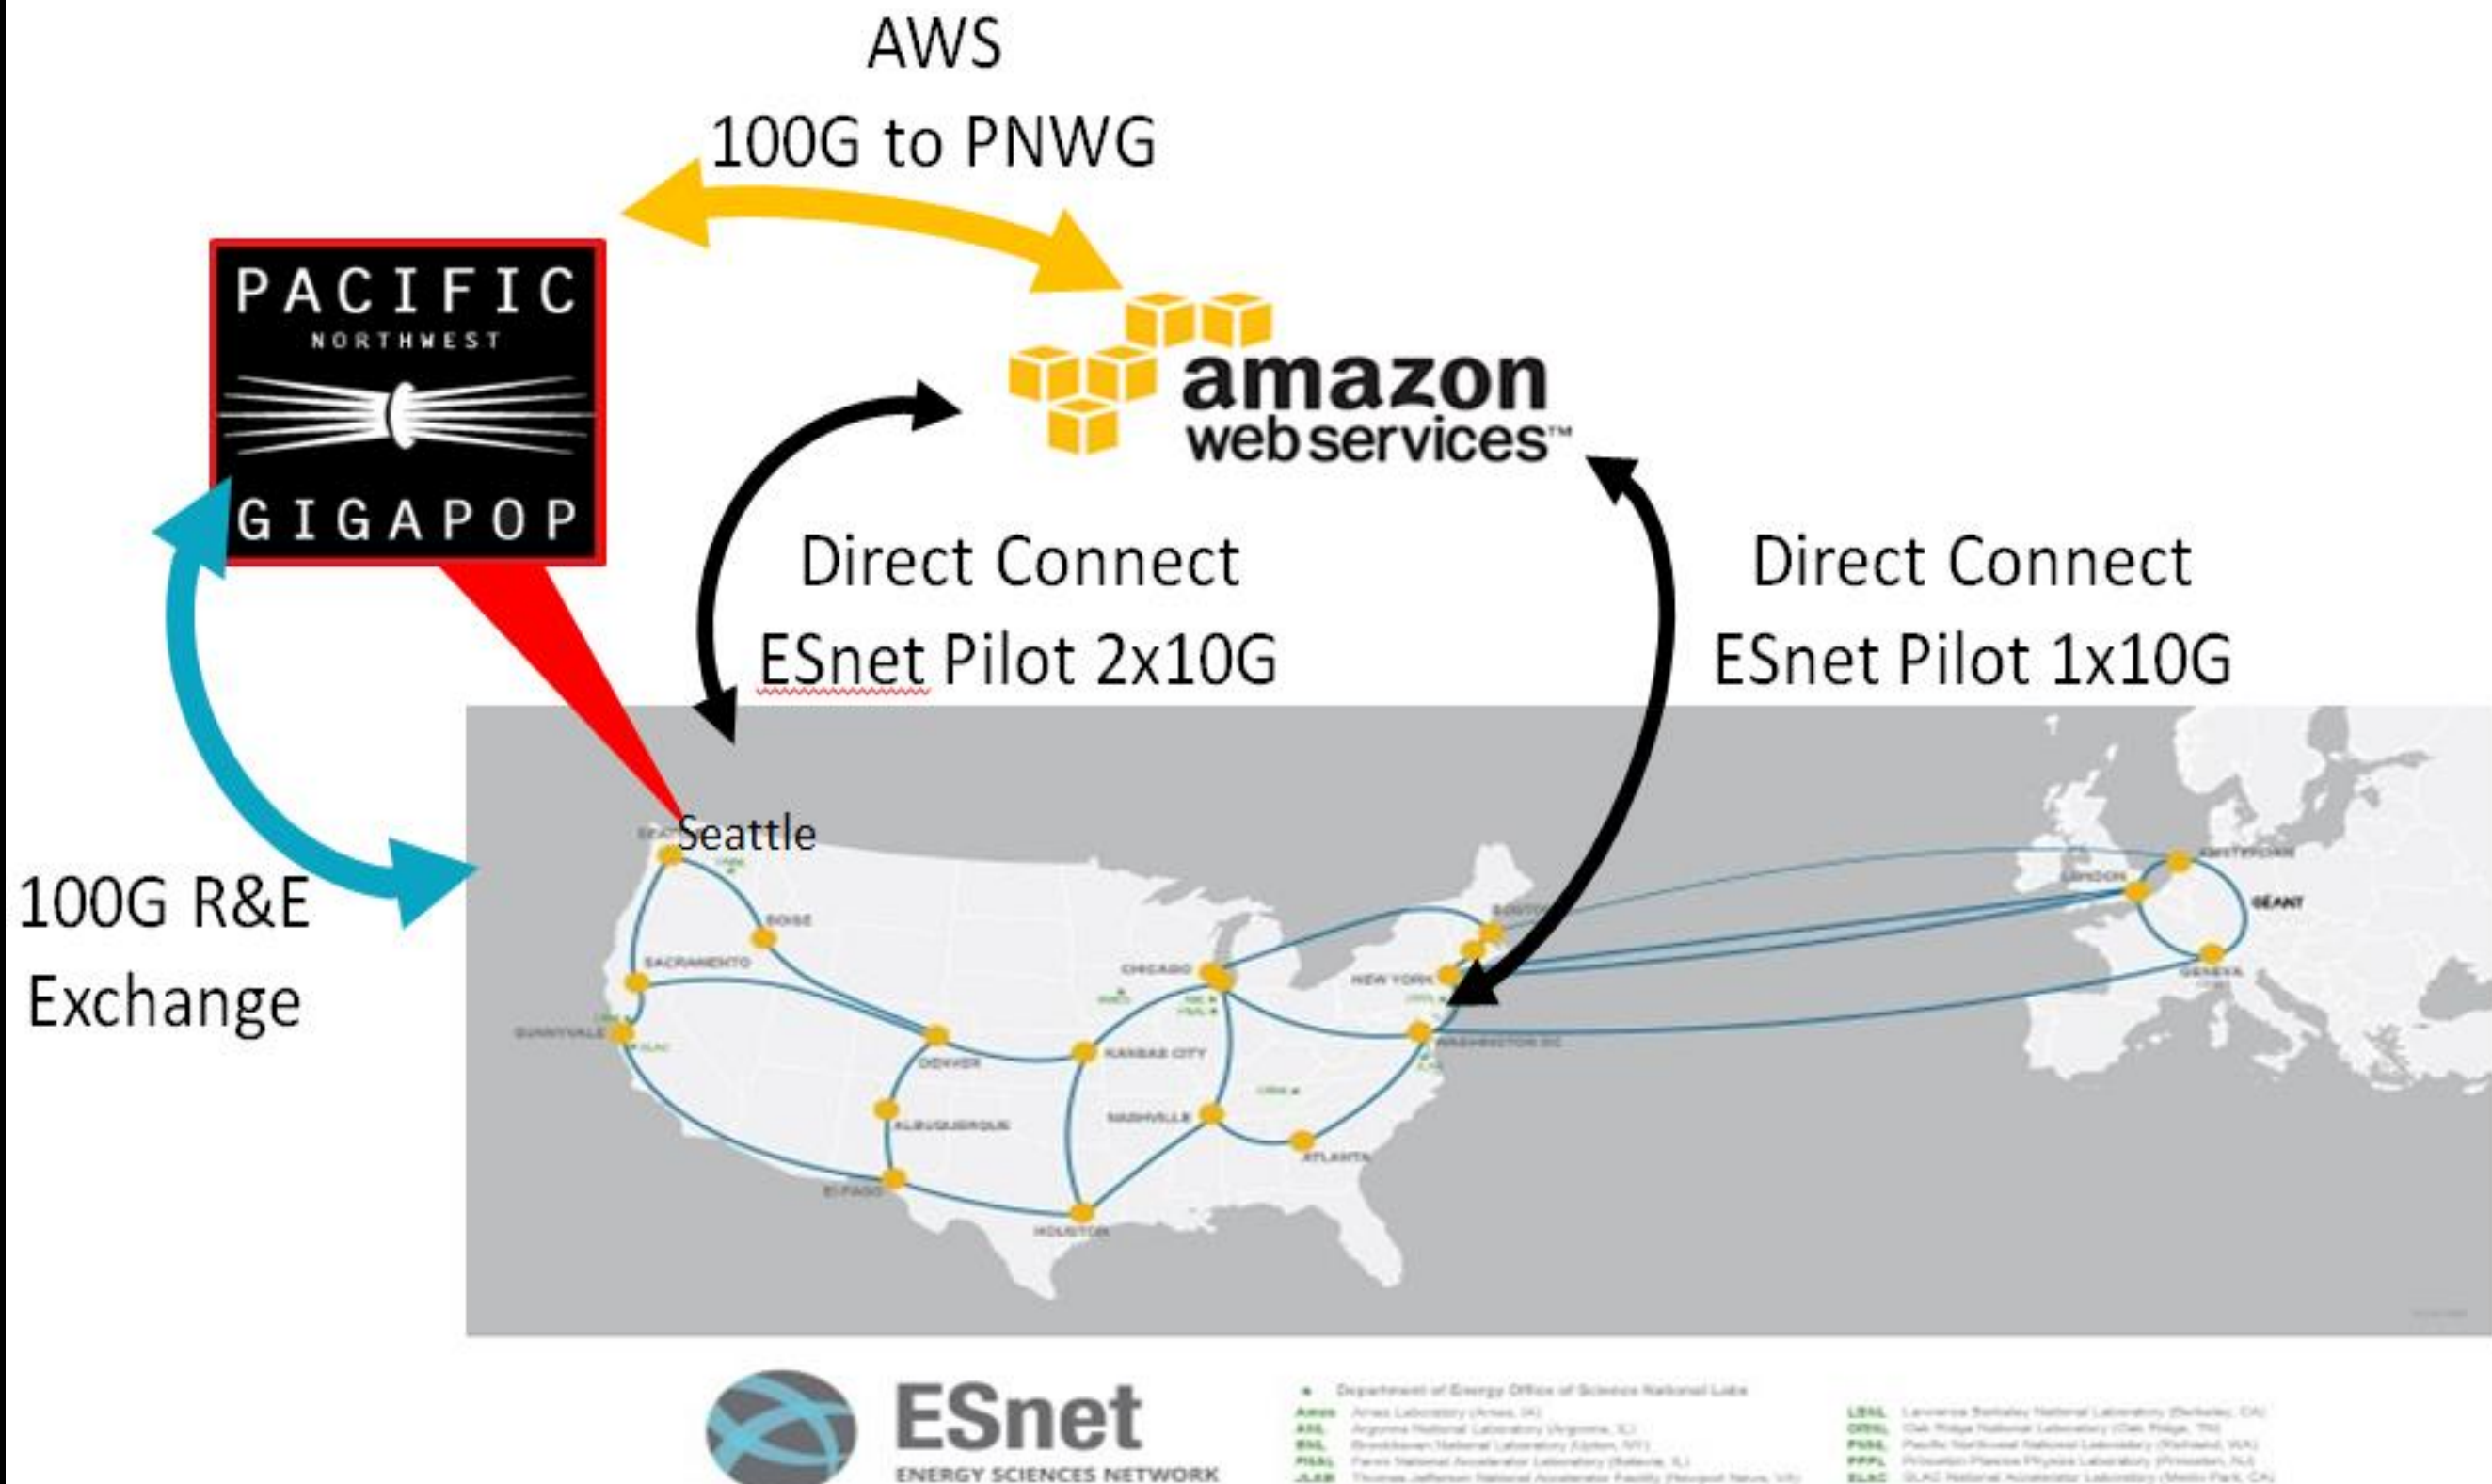
In a previous run we've observed 10-20% VM terminations (out of 2500 VM instances) with 2-hour jobs -> the shorter the job the better

- ATLAS needs to match volatile and opportunistic resources with workload profile that suits the characteristics of a volatile resource
 - Minimize loss due to resource becoming unavailable at any point in time
 - The Event Server comes with all features that perfectly fit the characteristics of the AWS spot market
 - Fine-grained processing at the event level – if we lose a VM we lose no more than a single event
 - Supports parallel processing at high degree – can “grab” and utilize as many CPU resources as the provider can offer
 - Utilizes high performance Object Store technology – this is the storage technology cloud providers have been focusing on
-

Event Server based Physics Event Simulation on AWS



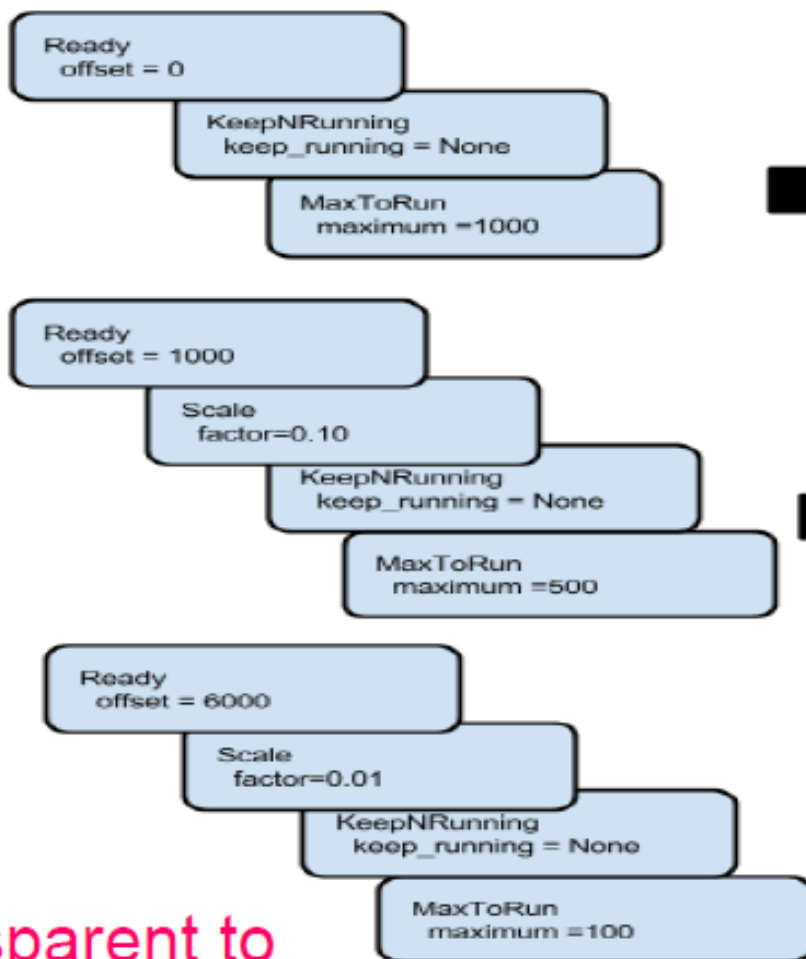
Connecting AWS Facilities to the Research Community



Using Cloud Resources effectively: A Policy-based Cloud Scheduler, developed at RACF

Example: Cascading Cloud Targets

Plugins and Config



Targets




Policy

Fill local Cloud 1-1 when up to 1000 jobs are waiting.

Trigger EC2 spot VMS 1 per 10 waiting jobs over 1000. Never run more than 500.

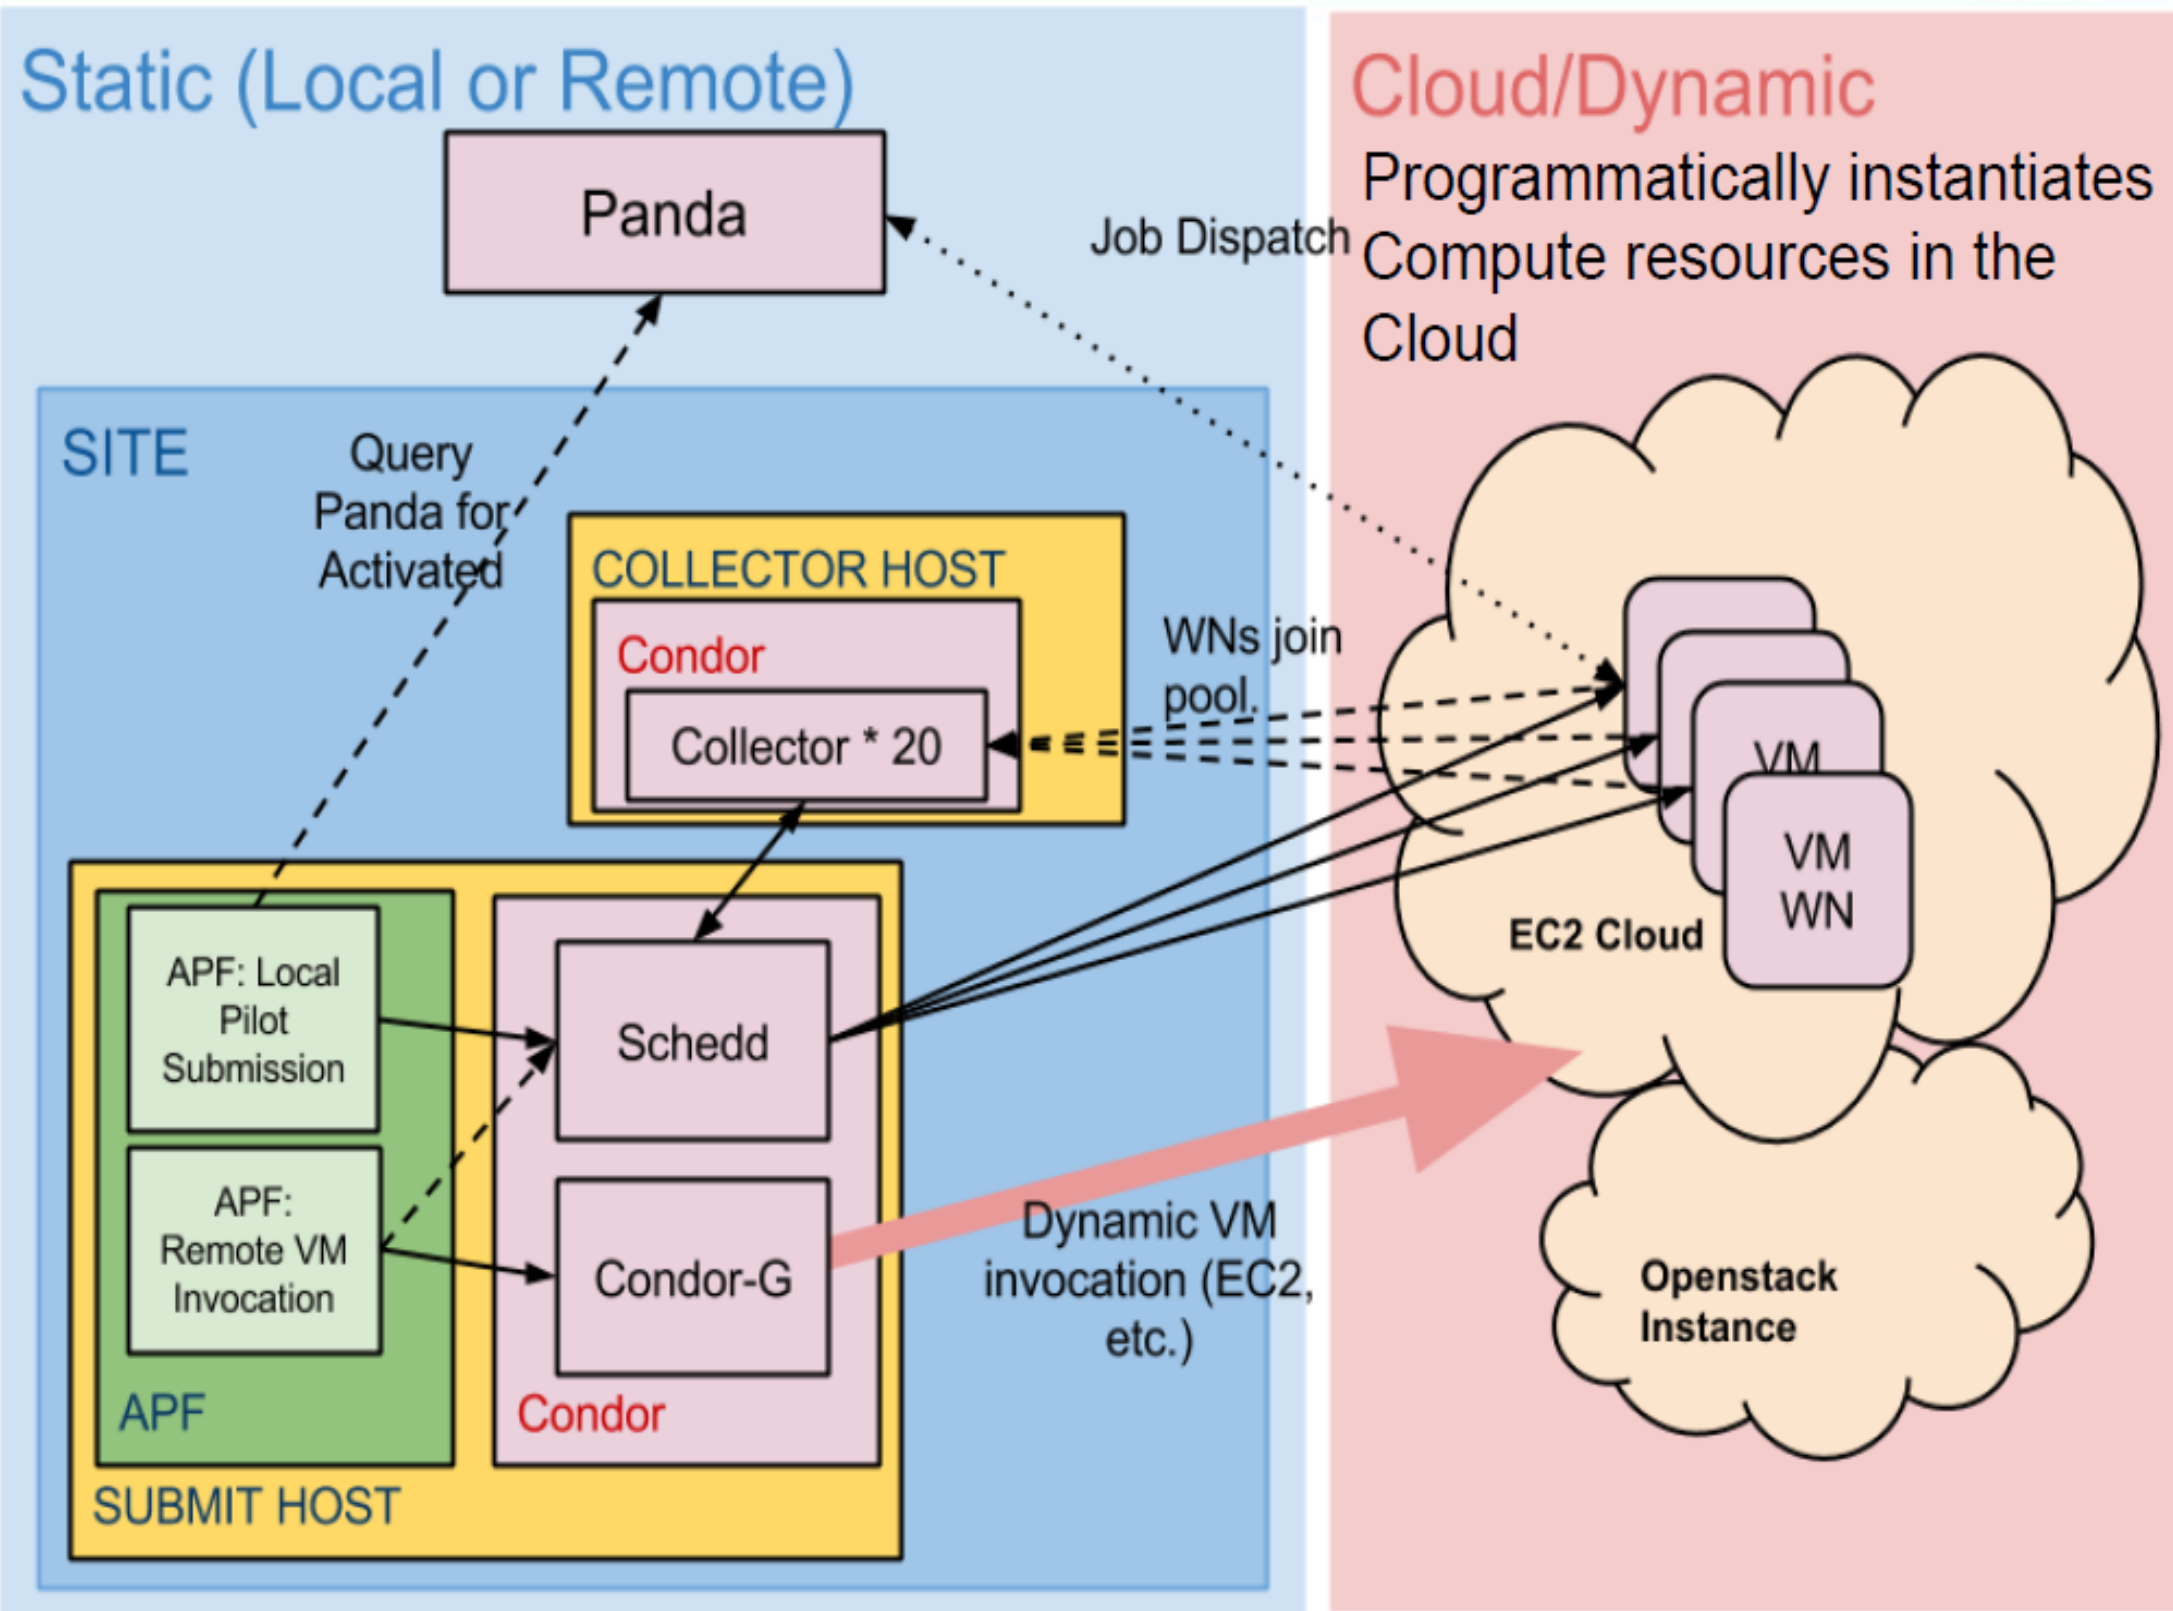
Trigger EC2 on-demand VMS 1 per 100 waiting jobs over 6000. Never run more than 100.

Fully transparent to Workload Management System (e.g. PanDA), Elastically expands Pool of Compute Resources according to user-defined policy




Demand-driven, policy-based programmatic instantiation and contraction of cloud resources

Elastic Cluster



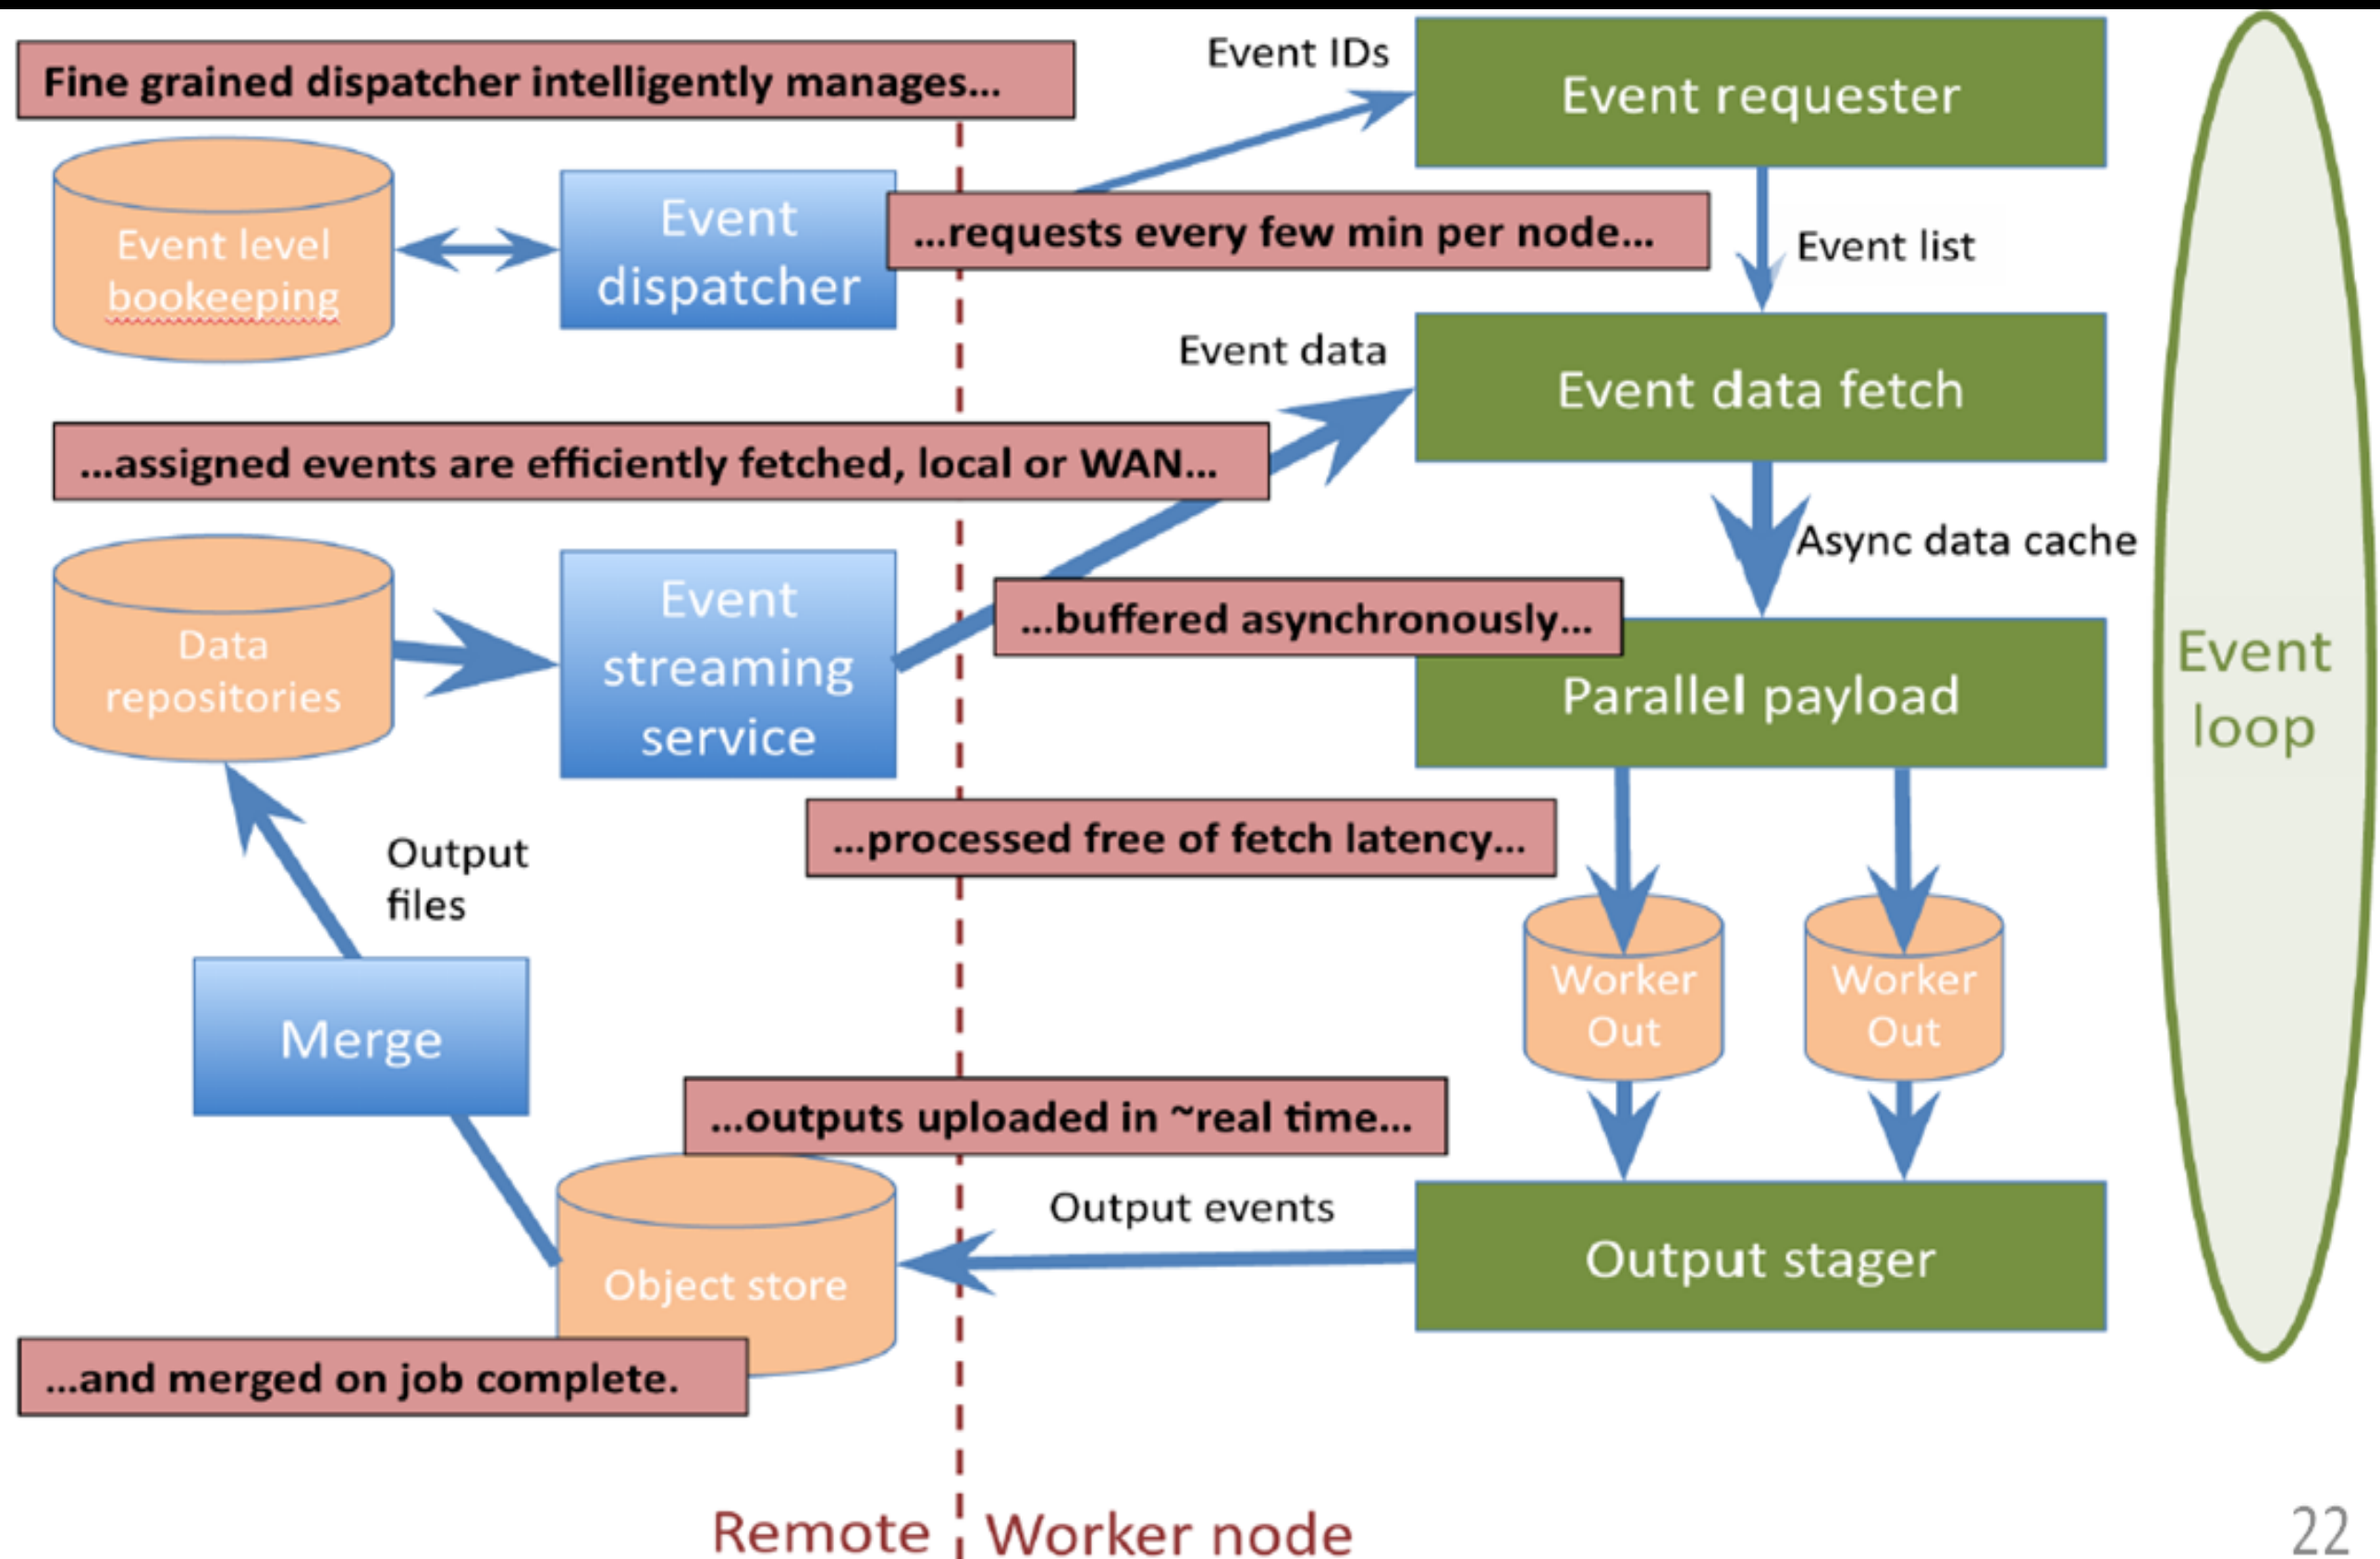
John Hover, BNL



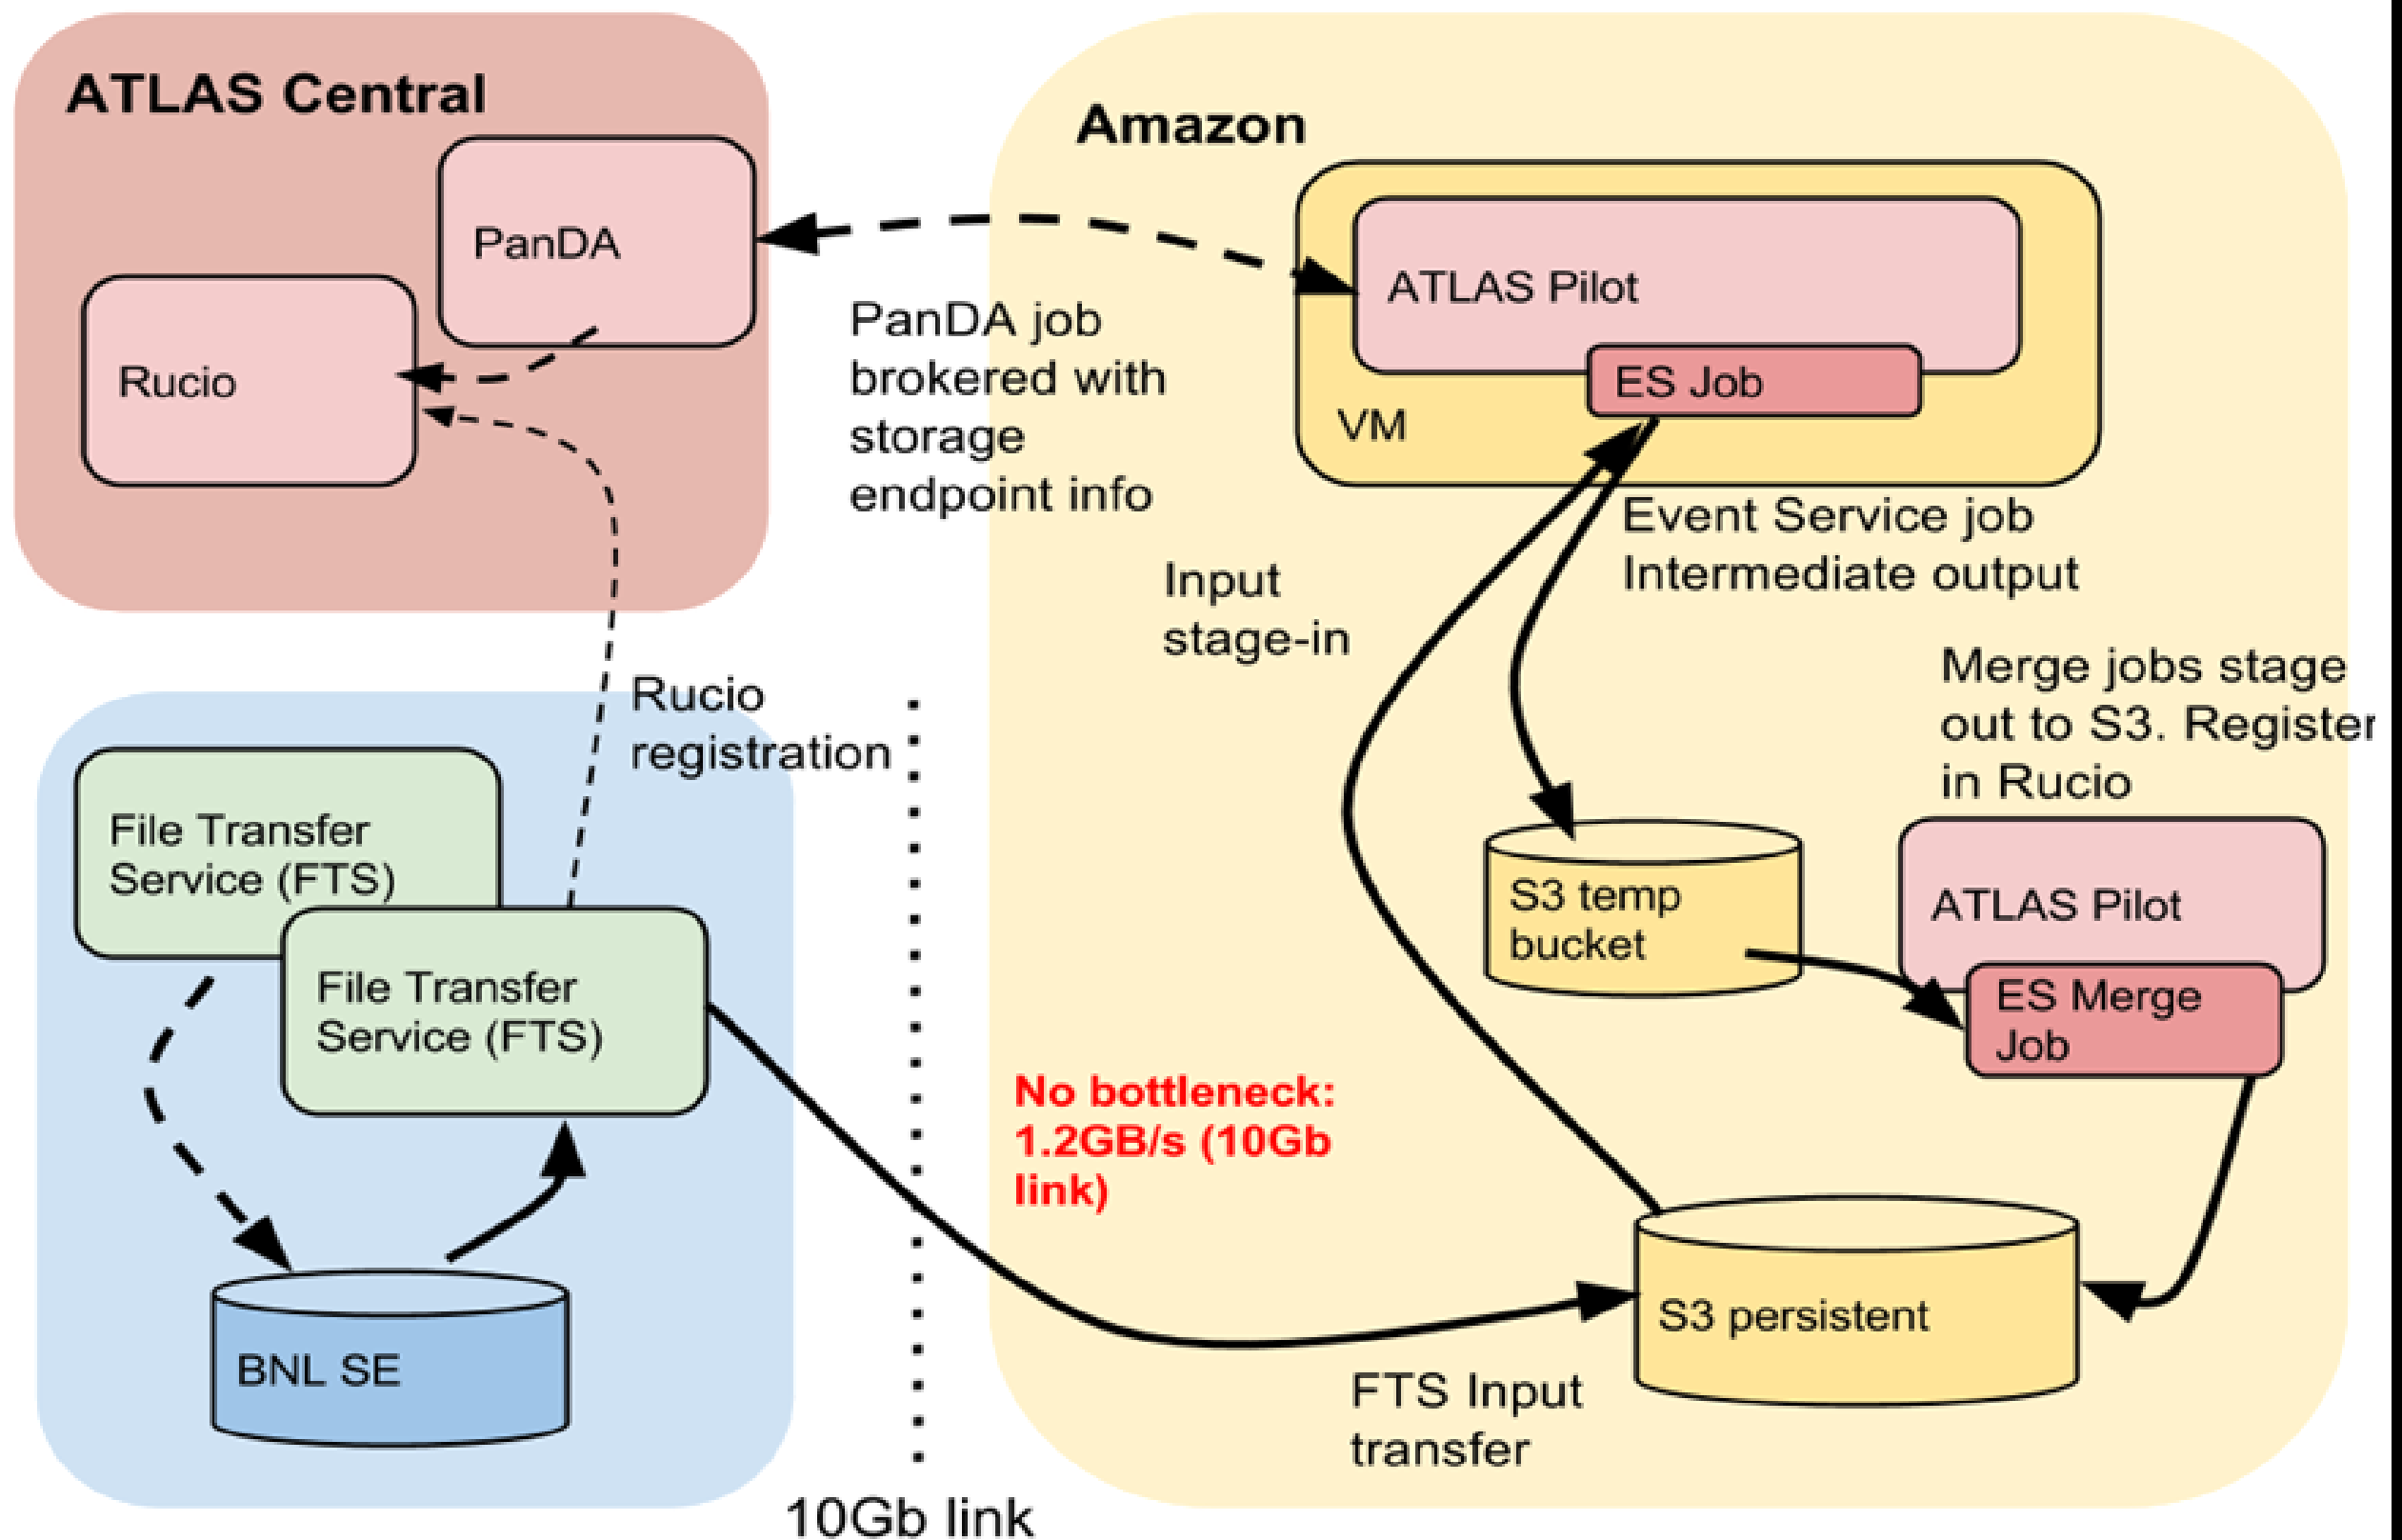
ATLAS Event Processing

- **ATHENA** – ATLAS simulation, reconstruction and data analysis
 - **Huge Codebase:** millions of lines of code organized into thousands of independent software packages
 - Designed and developed over many years for **sequential processing only**
 - Static workload
 - Job processes a **predefined sequence of events** from a given file
 - Only after **processing all events** from the statically allocated workload the job is considered
 - In case of **premature** termination the job is considered **failed** and **all data produced by the job is lost**

The ATLAS Event Service

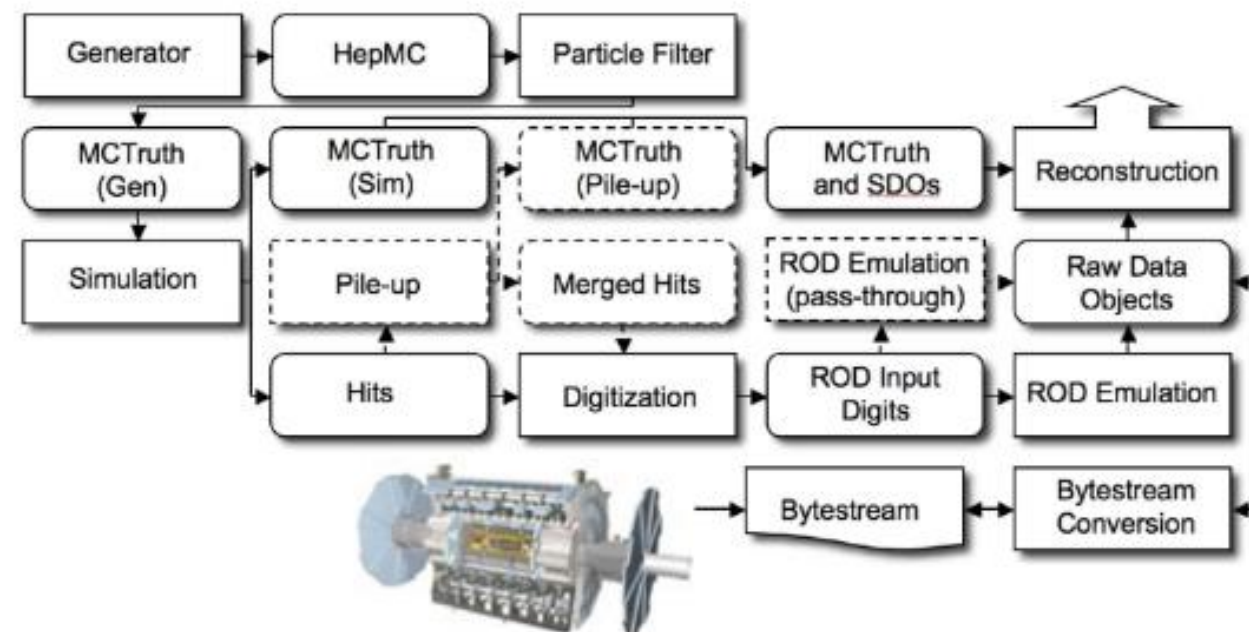



Architectural Overview from the Facility Perspective



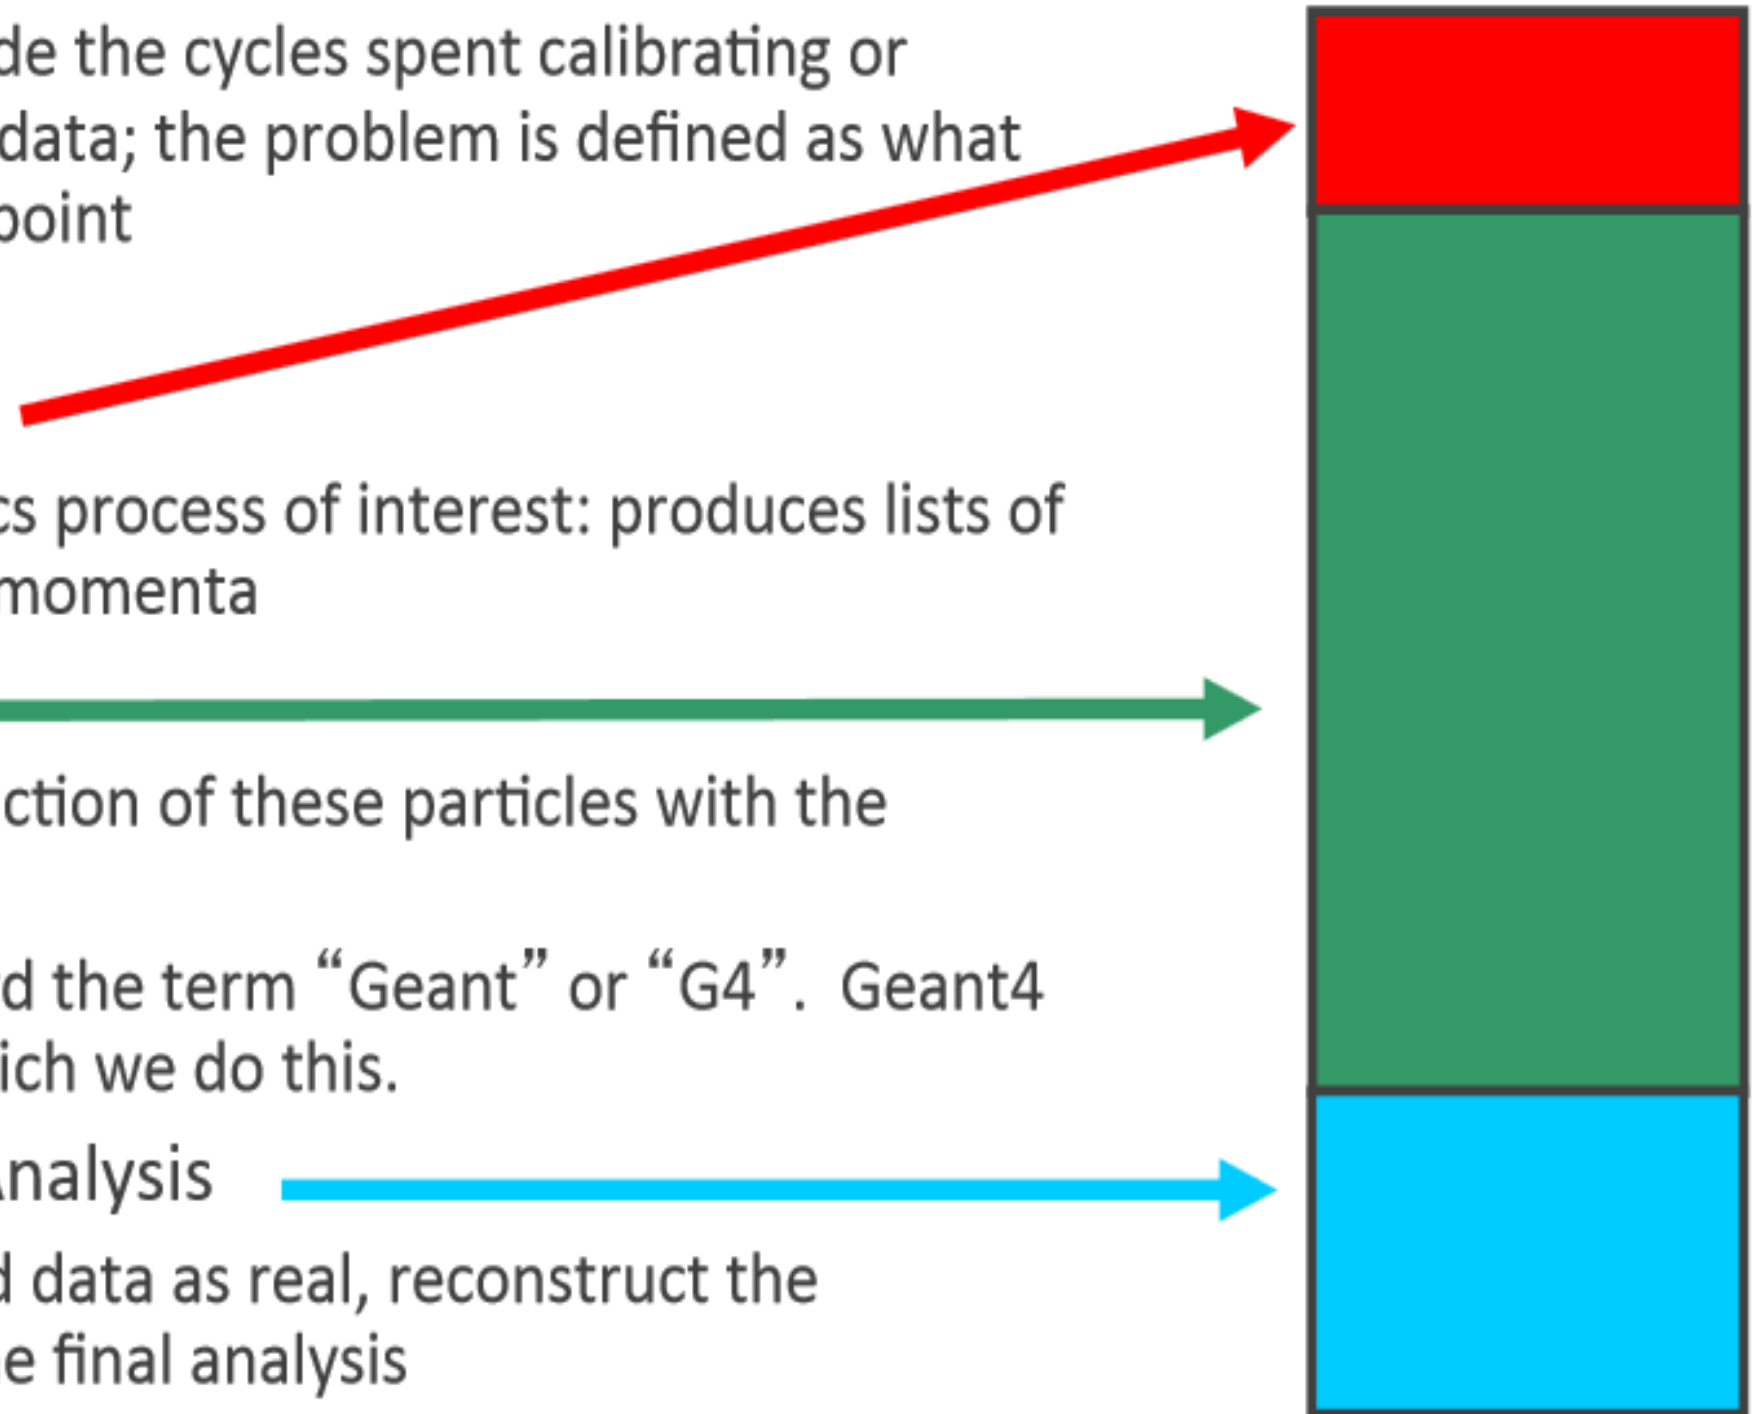
Simulation is important and resource intensive

- We collide particles together and measure the trajectories of the products in our detectors
- We then compare these results with simulation – at multiple levels
 - Does the detector respond to these particles in the way we expect?
 - Do we see the number of particles in various categories that we expect?
 - Etc.
- This is a complex process, and *we are as dependent on the simulation chain as we are on the data chain.*
 - ATLAS uses about a billion cpu-hours per year on this.



Computing to reach the Science Goals

- ATLAS uses about a billion CPU-hours per year on the Grid
 - This does not include the cycles spent calibrating or reconstructing the data; the problem is defined as what happens after this point
- Event Generation
 - Simulate the physics process of interest: produces lists of particles and their momenta
- Simulation
 - Simulate the interaction of these particles with the detector
 - You may have heard the term “Geant” or “G4”. Geant4 is the toolkit by which we do this.
- Reconstruction and Analysis
 - Treat the simulated data as real, reconstruct the particles, and do the final analysis



The Power of a Supercomputer

- We can run using the entire machine
 - For throughput reasons, we normally limit ourselves to 1/3 of the machine: a million parallel processes
- Per core event generation rate is ~1.5x a Grid core
 - Speedup of x23 over the year
 - Mira has a lot of cores! (768432)

Leadership Computing Facility

Mira Activity

	R00	R01	R02	R03	R04	R05	R06	R07	R08	R09	R0A	R0B	R0C	R0D	R0E	R0F
M1																
M0																
	R10	R11	R12	R13	R14	R15	R16	R17	R18	R19	R1A	R1B	R1C	R1D	R1E	R1F
M1																
M0																
	R20	R21	R22	R23	R24	R25	R26	R27	R28	R29	R2A	R2B	R2C	R2D	R2E	R2F
M1																
M0																

While this job was running, Mira was producing the equivalent computing as 5 or 6 ATLAS Grids.

On our best days, we provide the equivalent computing capacity of the whole ATLAS Grid.



#5



Anselm



Triolith



#6



Edison




Abel,
Abisko



Kurchatov



Archer



*ATLAS Grid would
be around #27
from Top100*



#2

Titan System (Cray XK7)			
Peak Performance	27.1 PF 18,688 compute nodes	24.5 PF GPU	2.6 PF CPU
System memory	710 TB total memory		
Interconnect	Gemini High Speed Interconnect	3D Torus	
Storage	Lustre Filesystem	32 PB	
Archive	High-Performance Storage System (HPSS)	29 PB	
I/O Nodes	512 Service and I/O nodes		

DLGF 20

The ATLAS collaboration have members with access to these machines and to many others...



SuperMUC



#7
Stampede

ASCR Computing Upgrades

System attributes	NERSC Now	OLCF Now	ALCF Now	NERSC Upgrade	OLCF Upgrade	ALCF Upgrades	
Name Planned Installation	Edison	TITAN	MIRA	Cori 2016	Summit 2017-2018	Theta 2016	Aurora 2018-2019
System peak (PF)	2.6	27	10	> 30	150 → 200	>8.5	180
Peak Power (MW)	2	9	4.8	< 3.7	10	1.7	13
Total system memory	357 TB	710TB	768TB	~1 PB DDR4 + High Bandwidth Memory (HBM)+1.5PB persistent memory	> 1.74 PB DDR4 + HBM + 2.8 PB persistent memory	>480 TB DDR4 + High Bandwidth Memory (HBM)	> 7 PB High Bandwidth On- Package Memory Local Memory and Persistent Memory
Node performance (TF)	0.460	1.452	0.204	> 3	> 40	> 3	> 17 times Mira
Node processors	Intel Ivy Bridge	AMD Opteron Nvidia Kepler	64-bit PowerPC A2	Intel Knights Landing many core CPUs Intel Haswell CPU in data partition	Multiple IBM Power9 CPUs & multiple Nvidia Volta GPUS	Intel Knights Landing Xeon Phi many core CPUs	Knights Hill Xeon Phi many core CPUs
System size (nodes)	5,600 nodes	18,688 nodes	49,152	9,300 nodes 1,900 nodes in data partition	~3,500 nodes	>2,500 nodes	>50,000 nodes
System Interconnect	Aries	Gemini	5D Torus	Aries	Dual Rail EDR- IB	Aries	2 nd Generation Intel Omni-Path Architecture
File System	7.6 PB 168 GB/s, Lustre®	32 PB 1 TB/s, Lustre®	26 PB 300 GB/s GPFS™	28 PB 744 GB/s Lustre®	120 PB 1 TB/s GPFS™	10PB, 210 GB/s Lustre initial	150 PB 1 TB/s Lustre®

Today ↔ Future

Phase Space of available & affordable Resources

10,000 feet overview

Grid

- Virtual Organizations (VOs) of users trusted by Grid sites
- VOs get allocations
→ **Pledges**
 - Unused allocations: opportunistic resources

Trust Federation

Cloud

- Community Clouds - Similar trust federation to Grids
- Commercial Clouds - **Pay-As-You-Go** model
 - Strongly accounted
 - Near-infinite capacity → **Elasticity**
 - Spot price market

Economic Model

HPC

- Researchers granted access to HPC installations
- Peer review committees award **Allocations**
 - Awards model designed for individual PIs rather than large collaborations

Grant Allocation

This new paradigm significantly changes the role of Facilities in providing end-to-end solutions to their customers

- Facilities can benefit from providing a much more “elastic” offering
- Moving away from stand-alone stove-piped facilities toward Labs providing leadership role in the computing **eco-system**
 - ✦ labs, leadership class and production class facilities, etc are part of and leaders in the eco system, that includes scientific and commercial providers
 - ✦ sites will find a large “market” for specific offerings: specialized architectures, archival capabilities, database services, data management solutions
 - ✦ the eco-system is enabled by the labs and OSG and others
- Facility’s role is still to provide “complete solutions” for their users
 - ✦ CPU and data capacities with guaranteed level of service
 - ✦ Users would not have to care about whether their jobs are running on “owned” or “rented” resources Sites could make the economic decision themselves and optimize their cost structure
 - ✦ Storage services that adapt to where the jobs are running
 - ✦ On-demand services that scale by tapping into large pooled resources
 - like clouds, HPC, OSG etc

Trends across the HPC Landscape

- **Increasing importance of computing and simulation within SC and DOE programs (across our missions: science, national security, and quest for cleaner energy)**
 - Continued full subscription of computational resources
 - Increasing importance of effective partnerships (domain + applied math + CS)
 - Difficulties persist in acquiring and retaining highly skilled workforce
- **Nexus of big data and powerful compute is an emerging frontier**
- **The drive toward exascale – Scientifically important and challenging questions await exascale**
 - Dennard scaling is driving further increases in concurrency → **billion-way concurrency is coming**
 - Post-CORAL computer architectures may be significantly different → **significant re-coding**
 - Data movement is increasingly costly → **becoming the rate limiter**
 - Power consumption remains an issue (\$1M/megawatt-year)
 - A petaflop in a 19-inch rack
- **The post-Moore's Law epoch is drawing nearer – we need to start preparing**

S. Binkley/ ASCR

Post Moore's Law Computing

- CMOS lithographic feature sizes are approaching fundamental limits
 - Currently at 22 nm (both Intel and Nvidia)
 - 11 nm is projected for ~~~2015~~ 2016 (both Intel and Nvidia)
 - However, gate lengths may be smaller than 6 nm – corresponding gate dielectric thickness may reach a monolayer or less
 - The industry roadmap reaches beyond 11 nm (7 nm and 5 nm) but may be unattainable
 - Non-silicon extensions of CMOS, e.g., using III-V materials or nanotubes/nanowires or non-CMOS technologies, including molecular electronics, spin-based computing, single-electron devices, and graphene have been proposed
 - At scales of ~10 nm, quantum tunneling may become significant
 - Capital costs for tooling are increasing dramatically as feature sizes shrink
- Options:
 - Computing using superconducting technologies
 - Quantum computing/quantum information science
 - Neuromorphic computing
 - Probabilistic computing
 - ???

**Considerable R&D
required**

Summary and Outlook

- There is no end in sight for large increases in resource demands and new capabilities, which change expectations and requirements on HEP facilities
 - ✦ provide services to distributed communities, supporting complex end-to-end use cases involving huge computational and data throughput needs and capabilities
- The role of the facility providers are changing as they are facing cost effective competition to their “bare metal” offerings from IaaS providers
 - ✦ Facilities remain to be first-line support for the complex scientific work flows and data management needs of HEP and other DOE SC communities
 - ✦ Facilities should integrate new opportunities and capabilities into their service offerings, in particular in connecting to large data management and data access systems, beyond “login and batch” services for applications and application libraries
- Facilities should keep an open mind how to provide their services so they fit into and enrich the US and international scientific computing eco-system
 - ✦ requires new thinking and approaches to difficult issues in the distributed environment, including security, robustness and protection of resources, accounting, prioritization etc
 - ✦ good experiences with LHC and emerging IF experiments
- HEP and other DOE Facilities clearly have a huge opportunity for great leadership roles in this environment